

Explainable AI Isn't Enough! Rethinking Algorithmic Contestability

Timo Freiesleben^{1,2,3}, Kristof Meding^{4,5}, Gunnar König^{4,6}

¹LMU Munich, Germany

²Munich Center for Machine Learning (MCML)

³Munich Center for Mathematical Philosophy (MCMP)

⁴University of Tübingen, Germany

⁵CZS Institute for Artificial Intelligence and Law

⁶Tübingen AI Center

timo.freiesleben@lmu.de

Abstract

Machine learning systems increasingly make life-changing decisions about individuals, such as loan approvals, hiring, and cheating detection, raising a pressing question: how can individuals respond to negative decisions made by these opaque systems? While explainable artificial intelligence (XAI) has largely focused on algorithmic recourse—helping individuals change their features to obtain a desired outcome—the parallel problem of algorithmic contestability—helping individuals review and correct erroneous algorithmic decisions—has received far less attention, despite its central ethical and legal importance. We trace this neglect to the absence of clear formal definitions and a systematic operationalization of contestability as an algorithmic problem. To address it, we propose an operational definition of contestability as a natural complement to recourse: contestability starts from the presumption that a decision may be incorrect and focuses on identifying evidence to challenge and potentially overturn it, whereas recourse assumes the decision is valid and instead provides pathways for changing it. We show that standard XAI explanations, such as counterfactuals, LIME, or Anchors, even when combined with human intuitions about decision continuity or monotonicity, reveal only errors in the neighborhood of the individual, but provide insufficient grounds for overturning the decision at hand. Going thus beyond traditional XAI, we identify three types of evidence warranting reversal according to the decision maker's own ethical standards: predictive multiplicity, incorrect feature values, and neglected overruling evidence. We argue that these render decisions normatively indefensible and thus successfully contestable. Finally, we analyze how existing EU legislation connects to our framework and argue that individuals already hold some legal rights to these forms of evidence.

1 Introduction

Machine learning systems increasingly make life-changing decisions about individuals. They influence whether laypeople receive loans, whether applicants are invited to job interviews, or even whether students are accused of cheating on an exam. Seeing one's life placed in the hands of an opaque algorithm can make people feel powerless: What can they do when confronted with an unfavorable decision?

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Explainable artificial intelligence has long been motivated, at least in part, by the aim of restoring the agency of individuals who face unfavorable algorithmic decisions (Lipton 2018; Wachter, Mittelstadt, and Russell 2017). Famously, *algorithmic recourse explanations* have been proposed to recommend actionable steps, e.g., paying off outstanding loans, that enable individuals to reverse an unfavorable decision. Another proposed way to restore individuals' agency is to provide them with grounds for *contesting* adverse decisions, that is, to provide individuals with information to review the decision-making process and intervene in situations where they feel judged incorrectly or unfairly. For example, individuals may contest a rejected credit application if their rejection was based on living in a neighborhood with high loan default rates; whether the contestation is successful and leads to a reversal of the decision is then subject to a negotiation process.

While the literature on algorithmic recourse has become a cornerstone of XAI research, with both solid conceptual as well as algorithmic foundations (Ustun, Spangher, and Liu 2019; Venkatasubramanian and Alfano 2020; König, Freiesleben, and Grosse-Wentrup 2023), the literature on algorithmic contestability has remained comparatively underexplored. This is striking given contestability arguably has at least as great ethical and legal importance as algorithmic recourse (Selbst and Barocas 2018; Lyons, Velloso, and Miller 2021; Bayamlioğlu 2022). The reasons for this disparity are undoubtedly multifaceted: contestability is deeply intertwined with legal questions that vary across jurisdictions; technical researchers may feel uncomfortable engaging with legal concepts; and decision makers often have little incentive to provide individuals with evidence that could be used to challenge their decisions. However, we believe there is a more fundamental reason why contestability has been largely neglected as an algorithmic problem: the absence of clear formal definitions and a systematic operationalization of contestability itself. The goal of this paper is to provide precisely these.

We conceptualize algorithmic contestability as the counterpart to algorithmic recourse, with both approaches ultimately aimed at helping individuals achieve their desired outcomes (Section 3). In the context of recourse, individuals trust the correctness of the model's decision and use recourse recommendations to obtain the desired outcome.

That is, they believe that the decision about the individual (e.g., loan rejection) aligns with the true state (i.e., the loan would not be repaid) and the recourse explanation provides information that enables them to change both the decision and the true state, that is, the loan is approved and ultimately repaid. Algorithmic contestability takes the opposite perspective: the individual maintains that the decision (e.g., loan rejection) is not aligned with the true state (i.e., the loan would have been repaid) and therefore seeks information that allows them to alter the decision to match the true state, so that they receive the loan.

We show that common XAI methods, such as counterfactual explanations (Wachter, Mittelstadt, and Russell 2017), Local Interpretable Model-Agnostic Explanations (LIME) (Ribeiro, Singh, and Guestrin 2016), and Anchors (Ribeiro, Singh, and Guestrin 2018), are of limited value for algorithmic contestability. While they can reveal conflicts with human intuitions concerning continuity and monotonicity, and thereby point to errors somewhere in the individual’s neighborhood, they do not provide sufficient grounds to overturn a decision (Section 4). As a response, we move beyond XAI and investigate other types of evidence that can enable contestability from a *normative perspective*, that is, evidence that should lead decision makers to change their decision according to their own ethical norms (Section 5). The three types of evidence are 1. *Predictive multiplicity*: conflicting predictions for the individual from a set of models with equal empirical and theoretical support; 2. *Incorrect feature values*: wrong information about the individual that if corrected would lead to a different decision; 3. *Neglected overruling evidence*: evidence about other features of the individual that if taken into account would overturn the decision. We show that these forms of evidence can be sufficient to alter algorithmic decisions on normative grounds. We conclude by linking the European General Data Protection Regulation (GDPR) to our framework and argue that individuals could benefit from rights to these forms of evidence under the existing legislative landscape (Section 6).

2 Related Work

Contestability is a core ethical requirement in algorithmic decision-making and, consequently, a central element in several worldwide legislations (Vaccaro et al. 2019; Maxwell and Dumas 2023). For example, it is one of the eight guiding principles of Australia’s AI Ethics Framework (Department of Industry, Innovation and Science 2019), lies at the “heart of legal rights that afford individuals access to personal data and insight into decision-making processes” (Mulligan, Klutetz, and Kohli 2019), and constitutes the “backbone” of the widely discussed Article 22 of the GDPR (Bayamhoğlu 2022). The goals and nature of contestability are, however, multifaceted and often debated, which makes both its legal and technical implementation challenging (Lyons, Velloso, and Miller 2021). It is widely agreed that algorithmic contestability remains under-researched, especially from a technical perspective (Lyons, Velloso, and Miller 2021; Alfrink et al. 2023; Maxwell and Dumas 2023; Yurrita et al. 2025).

Henin and Le Métayer (2021) provide a norm-based formal framework of contestability that allows decisions to be

contested by reference to norms and justifications. Their formal model, *Algocate*, provides an interactive tool for generating evidence in favor of and against a decision. While their idea of building upon norms and justifications is conceptually sound (Aler Tubella et al. 2020), their technical framework is difficult to implement in practice: it requires a complete logical specification of decision norms as well as an ordering relation over argument strengths. More recently, Moreira et al. (2025) attempted to formalize contestability in a comparatively broad sense, encompassing the facilitation of risk analysis, the assurance of accountability, and the enablement of decision reversal. Unlike this work, we focus exclusively on contestability to enable decision reversal and provide a concrete framework together with workable definitions for this specific scenario. Since we narrow our focus to enabling decision reversal through contestability, we abstract away from other aspects in the construction of algorithmic systems that may be contestable (Hirsch et al. 2017; Sarra 2020; Alfrink et al. 2023; Yurrita et al. 2025).

While providing individuals with grounds to contest individual decisions is often cited as a central motivation of the field of explainable artificial intelligence (XAI) (Ribeiro, Singh, and Guestrin 2016; Lipton 2018; Adadi and Berrada 2018; Molnar 2020), only very few technical works address contestability in greater detail (Mansi, Karusala, and Riedl 2025). One notable exception is the work on counterfactual explanations by Wachter, Mittelstadt, and Russell (2017). However, while Wachter, Mittelstadt, and Russell (2017) motivate *why* counterfactuals can enable individuals to contest decisions and satisfy the legal requirements established by the GDPR, they do not demonstrate *how* their method can be tailored to this purpose, nor under which conditions contestation based on counterfactuals should be successful. Other explanation methods, such as LIME or Shapley values, are also commonly discussed as potentially fostering contestability (Mulligan, Klutetz, and Kohli 2019; Yurrita et al. 2023). Many ethical and legal scholars argue that XAI could serve as an “enabler for contestability” (Selbst and Barocas 2018; Lyons, Velloso, and Miller 2021; Vredenburg 2022; Maxwell and Dumas 2023; Schmude et al. 2025). Yet, they often remain vague about *how* exactly XAI facilitates such contestation. The most detailed account is offered by Selbst and Barocas (2018), who argue that explanations must be paired with human intuitions to support contestability. However, as we will demonstrate, the insights that common XAI methods provide into the (in)correctness of decisions are quite limited, even when individuals actively draw on their own intuitions.

A closely related problem to algorithmic contestability is the problem of algorithmic recourse (Ustun, Spangher, and Liu 2019). Both algorithmic recourse and contestability aim to reverse decisions; however, in recourse, the individual perceives the decision as correct, whereas in contestability, it is perceived as incorrect (Venkatasubramanian and Alfano 2020). The recourse literature has evolved significantly over the years (Karimi et al. 2022), with work investigating the robustness (Upadhyay, Joshi, and Lakkaraju 2021; König et al. 2025), fairness (Von Kügelgen et al. 2022), and privacy risks (Pawelczyk, Lakkaraju, and Neel 2023) of recourse.

We hope that our definitions and operationalizations of algorithmic contestability will enable a similar formal treatment of contestability, allowing many ideas and methods to be transferred and ultimately placing algorithmic contestability on equal footing with algorithmic recourse.

3 From Algorithmic Decision-Making to Algorithmic Contestability

Algorithmic decision-making is characterized by a few core components. Individuals possess *unobserved properties* that decision makers would like to know, but that are difficult or sometimes impossible to measure directly. For example, in lending, credit institutions want to know whether individuals will ultimately repay their loans, and in cheating detection, schools want to know whether a student has cheated. These unobserved properties lie at the heart of decision makers’ *core ethical principles*¹ that guide their decisions. Ethical does not mean that these principles are justified morally but only that they are maxims the decision maker commits to. For instance, banks aim to grant loans only to individuals who will be able to repay them, and schools aim to fail students who have cheated. Ideally, decision makers would make decisions strictly in accordance with these core ethical principles; we call such decisions *normatively correct*.

Definition 1. Let i be an individual with an unobserved binary attribute $y^{(i)} \in \{0, 1\}$. We define the *normatively correct* decision as $d^*(i) := y^{(i)}$.²

However, because these properties are unobserved, decision makers must infer them from *features*, i.e., measured properties that are highly correlated with the unobserved property of interest. Examples include age, salary, or credit utilization in lending, or a student’s submitted answers in cheating detection. These features allow decision makers to predict the unobserved property with some probability, but such predictions are never perfect. For instance, although a young person with a high salary and low credit utilization is very likely to repay her loan, she may still default due to unforeseen circumstances such as job loss or an accident. Decisions based on such predictions therefore require decision makers to accept a certain *inductive risk* (Douglas 2000). For example, they may choose to reject loan applicants whose probability of repayment falls below a given threshold. Accepting inductive risk entails accepting the possibility of normatively incorrect decisions: some applicants will be approved despite ultimately defaulting, and some students will be accused of cheating despite having done nothing wrong. Setting the level of inductive risk is, ultimately, an ethical decision that decision makers must make (Douglas 2009). We call decisions that adhere to the chosen inductive risk *epistemically correct*.

¹To support readers with more technical backgrounds, the appendix of the extended version of this paper <https://arxiv.org/abs/2605.16041> offers a short glossary explaining the philosophical concepts used throughout the paper.

²While standard in machine learning theory, the assumption that such ground-truth labels exist for unobserved attributes is controversial in many contexts (Cabitza, Campagner, and Basile 2023; Kang 2023).

Definition 2. Let $\tau \in [0, 1]$ denote the inductive-risk threshold, and let $p(x^{(i)}) := \mathbb{P}(Y = 1 \mid X = x^{(i)})$ be the probability that individual i possesses the unobserved property $Y = 1$ given her features $x^{(i)}$. We define the *epistemically correct* decision given features X as

$$d(x^{(i)}) := \begin{cases} 1 & \text{if } p(x^{(i)}) > \tau, \\ 0 & \text{else.} \end{cases}$$

Decision makers do not have access to the true probabilities of loan default or cheating conditional on feature values. Instead, they must estimate these probabilities from data, typically using machine learning methods. The machine learning model selection process involves a variety of choices—concerning model classes, regularization techniques, hyperparameters, and more (this pipeline is discussed in more detail in Section 5.1)—with the goal of training a model on a dataset that achieves high predictive performance. Moreover, decision makers usually do not have access to individuals’ true feature values. In lending, for example, individuals may provide incorrect information, or lenders may store inaccurate data.

Definition 3. Let $\hat{p}$ denote the model resulting from this model selection process, and let $\tilde{x}^{(i)}$ denote the feature values available to the decision maker. We define the *actual decision* as

$$\hat{d}(\tilde{x}^{(i)}) := \begin{cases} 1 & \text{if } \hat{p}(\tilde{x}^{(i)}) > \tau, \\ 0 & \text{else.} \end{cases}$$

Algorithmic *decisions* thus result from two components: decision makers have ethical norms that determine which decisions are (normatively or epistemically) correct according to their core ethical principles and chosen inductive risk, and these ethical norms yield actual decisions only through their combination with epistemic beliefs. In the context of machine learning, these epistemic beliefs depend on measured feature values and on the prediction model induced by a learning algorithm.

That a decision is normatively or epistemically correct in our terminology does not mean that it is uncontroversial, only that it aligns with the decision maker’s ethical norms. One may, for instance, contest a school’s core ethical principle and argue that using generative AI in exams should not be penalized (Anders 2023). Likewise, decisions that are epistemically correct can still be challenged on ethical grounds, for example, by arguing that a threshold-based policy in cheating detection involves an unacceptably high inductive risk of false accusations. In this work, however, we focus not on challenging ethical norms but on challenging the epistemic beliefs underlying algorithmic decisions. Which decisions should be contestable on such grounds?

Definition 4. Assume a decision $\hat{d}(\tilde{x}^{(i)})$ is made about an individual i . We call the decision

- *normatively contestable* if the decision is normatively incorrect, i.e., $\hat{d}(\tilde{x}^{(i)}) \neq d^*(i)$, and
- *epistemically contestable* relative to a feature set X if the decision is epistemically incorrect, i.e., $\hat{d}(\tilde{x}^{(i)}) \neq d(x^{(i)})$.

	Level	Meaning	Evidence	↑ stronger contesting grounds
	Normative contestability	The decision $\hat{d}(x^{(i)})$ misjudges the actual label $y^{(i)}$.	Overruling evidence (Section 5.3)	
	Epistemic contestability	The decision $\hat{d}(x^{(i)})$ misjudges what the features warrant $d(x^{(i)})$.	Incorrect feature values (Section 5.2), Predictive multiplicity (Section 5.1)	
	Somewhere contestability	Some nearby x is epistemically contestable.	Reason intuitions + Anchors (Section 4.3), Continuity intuitions + Counterfactuals (Section 4.1)	
	Somewhere inaccuracy	The probability estimate $\hat{p}$ is off somewhere nearby.	Monotonicity intuitions + LIME (Section 4.2)	

Table 1: **Levels of contestability.** The table summarizes our main contributions in this paper: 1) We introduce four levels of algorithmic contestability, namely normative, epistemic, somewhere contestability, and local inaccuracies; 2) We show that XAI explanations, even when combined with human intuitions, only allow one to reach the lowest two levels; 3) We show that other types of evidence, such as predictive multiplicity, support higher levels of contestability. The arrow on the side indicates that establishing higher levels of contestability provides stronger normative grounds for overturning a decision.

Normative contestability binds decision makers to their own ethical principles. For example, if a student is accused of cheating but did not in fact cheat, she should be able to contest and overturn the decision as a matter of moral principle. Epistemic contestability, by contrast, binds decision makers to the inductive risk they have accepted. For instance, if a decision maker assigns a high default risk to an individual who is in fact low risk, the individual should be able to challenge and reverse the decision on epistemic grounds.

Anchoring contestability in decision makers’ own principles has a key advantage: most disputes over algorithmic decisions are unlikely to reach the courts and are instead resolved through negotiation between the individual and the decision maker. If decision makers remain committed to their ethical and epistemic principles, then decisions that are normatively or epistemically contestable should, provided individuals can present sufficient supporting evidence, be overturned.

4 XAI Explanations Can Only Establish That Errors Occur Somewhere

Model-agnostic XAI techniques provide insight into a machine learning model by analyzing its sensitivity to changes in the input (Scholbeck et al. 2019). Prominent examples include counterfactual explanations, LIME, or Anchors.

Why do scholars regard XAI as a key enabler of algorithmic contestability? This belief is rooted in an idea that has been articulated by Selbst and Barocas (2018) and is illustrated in Figure 1: suppose an explanation method, such as LIME, indicates that the requested loan amount has a locally positive effect on the model’s prediction, implying that requesting a higher loan increases the likelihood of approval. An individual may find this relationship highly counterintuitive, expecting instead that requesting a larger loan would reduce her chances of obtaining it. In this way, the combination of *human intuitions* and *model explanations* can reveal potential model errors and thereby seem to provide grounds for contestation. In the following, we show that while this

line of reasoning is not misguided, the grounds for contestation it yields are generally weak.

But before we do so, we must clarify what we mean by *human intuitions*. Rather than estimating conditional probabilities over unobserved variables, as machine learning models do, humans typically rely on *rules of thumb* when forming beliefs and making decisions (Gigerenzer and Gaissmaier 2011; Selbst and Barocas 2018). These rules are not arbitrary; they often exhibit high validity, having emerged through repeated experience and pattern recognition (Gigerenzer 2023). Both experts and non-experts draw on a rich repertoire of such intuitive rules, including the following: if a person highly similar to me obtains the loan, then so should I (*continuity rule*); a person with a higher income than me should, all else equal, have a higher probability of obtaining the loan (*monotonicity rule*); and certain properties of mine constitute sufficient or insufficient reasons for granting or denying the loan (*reasons rule*). We call a human intuition *correct* if the rule it expresses applies in the given context.

Suppose that an XAI explanation conflicts with human intuitions, as in the example above. To better understand what can be inferred from such a conflict, we introduce two notions of error. *Somewhere contestability* refers to the existence of epistemically incorrect decisions within a neighborhood of the individual, whereas *somewhere inaccuracy* refers to the existence of incorrect model predictions in such a neighborhood. More formally:

Definition 5. Let i be an individual and $\mathcal{B}_\epsilon(\tilde{x}^{(i)}) := \{z \in \mathcal{X} \mid \|z - \tilde{x}^{(i)}\| < \epsilon\}$ be an ϵ -ball around its feature values for a given $\epsilon > 0$. We call a decision *somewhere contestable* if there exists a $x \in \mathcal{B}_\epsilon(\tilde{x}^{(i)})$ such that $\hat{d}(x) \neq d(x)$. We call a machine learning model $\hat{p}$ *somewhere inaccurate* if there exists an $x \in \mathcal{B}_\epsilon(\tilde{x}^{(i)})$ such that $\hat{p}(x) \neq p(x)$.

Unlike epistemic and normative contestability, somewhere contestability and somewhere inaccuracies are insufficient to overturn a decision. The presence of an epistemically incorrect prediction somewhere in the neighborhood




Figure 1: **XAI explanations may conflict with human intuitions.** For example, (1) a counterfactual explanation may conflict with continuity rules, (2) a LIME explanation may conflict with monotonicity rules, and (3) anchors may contradict reason rules. Even under the strong assumption that these intuitions are correct, i.e. the rules apply, such conflicts imply an error only somewhere in the neighborhood of the explained observation and thus insufficient grounds to overturn the decision at hand. Details on the examples are reported in the Appendix of the extended version of this paper <https://arxiv.org/abs/2605.16041>.

of an individual remains fully compatible with the decision for that individual being correct. Logically, this yields a hierarchy: epistemic contestability is stronger than somewhere contestability, which in turn is stronger than somewhere inaccuracy.³ Table 1 provides an overview of all the different levels of contestability. Although errors somewhere in the neighborhood are a necessary condition for epistemically incorrect decisions and may thus serve as a starting point for algorithmic contestation, they provide only a very weak signal: most predictive models exhibit errors somewhere, leaving it largely indeterminate whether the specific decision at hand ought to be overturned.

In the following, we show that a conflict between human intuitions and XAI explanations—even in the best-case scenario where those intuitions are correct—supports at most an inference to an error somewhere. When the intuitions are themselves incorrect, which happens regularly (Tversky and Kahneman 1974), XAI explanations do not support any inference about the correctness of the decision of interest.

4.1 Detecting Continuity Conflicts With Counterfactual Explanations

Continuity rules take the form: *If two individuals have highly similar feature values, they should be judged similarly.* Such rules are widely adopted, including in the context of fairness (Kusner et al. 2017), as humans generally expect decisions to be made consistently (Kolodner 2014; Selbst and Barocas 2018).

Definition 6 (Continuity rule). If person i and j are very similar in their measured features $\tilde{x}^{(i)}$, they should receive the same decision, i.e. $d(\tilde{x}^{(i)}) = d(\tilde{x}^{(j)})$.

However, discontinuities can be justified. Consider two highly similar loan applicants, where one is 17 and the other is 18 years old. Granting the loan to the latter but not the former is justified, as they differ in legal capability, which is relevant in lending contexts.

³We provide a short proof for this claim in the appendix of the extended version of this paper <https://arxiv.org/abs/2605.16041>.

To check whether algorithmic decision-making conflicts with continuity rules, individuals need access to the feature values of similar alternative individuals who receive a different decision. We believe that counterfactual explanations are the best available approach to generate and select such alternative feature values.⁴ A counterfactual identifies feature values that are maximally similar to those of the individual but would have resulted in a different decision.

Definition 7 (Counterfactual explanation). Let $x \in \mathcal{X}$, then a counterfactual x_c is an alternative point in $\mathcal{X}$ that lies (maximally)⁵ close to x according to the distance function $\|\cdot\|$ but receives a different decision, that is, $\hat{d}(x) = 1$ and $\hat{d}(x_c) = 0$ or vice versa.

The following theorem (all proofs can be found in the appendix of the extended version of this paper <https://arxiv.org/abs/2605.16041>) shows that counterfactuals can provide sufficient information for somewhere contestability, however, as the counterexample (all counterexamples can be found in the appendix of the extended version of this paper <https://arxiv.org/abs/2605.16041>) shows, they do not render the decision in question epistemically contestable.

Theorem 1. *If the continuity rule applies to the individual and her counterfactual, and their distance is smaller than ϵ according to $\|\cdot\|$, the decision is somewhere contestable.*

⁴Closest real instances receiving a different classification would also be a reasonable choice, but raise data privacy concerns (Pawelczyk, Lakkaraju, and Neel 2023).

⁵While the original formulation of counterfactuals by Wachter, Mittelstadt, and Russell (2017) requires the counterfactual to be the closest input across the decision boundary, most methods generate counterfactuals heuristically, e.g. using evolutionary algorithms (Dandl et al. 2020), and therefore produce close rather than closest counterfactuals. For purposes of contestability, maximal proximity is less essential than it is for recourse (Venkatasubramanian and Alfano 2020; Freiesleben 2022).

4.2 Detecting Monotonicity Conflicts With LIME

Monotonicity rules take the form: *If an individual differs from another individual only on a feature that has a positive (or negative) relationship with the unobserved target, the former should not be judged worse or respectively better than the latter.* Such rules are widely adopted, as humans expect decisions to be affected by feature changes in a consistent direction (Birnbbaum 2019; Selbst and Barocas 2018).

Definition 8 (Monotonicity rule). If an individual i differs from an individual j only in feature k with $\tilde{x}_k^i \geq \tilde{x}_k^j$, then person i should receive at least as favorable a probability as j , i.e. $p(\tilde{x}^{(i)}) \geq p(\tilde{x}^{(j)})$ (respectively unfavorable $p(\tilde{x}^{(i)}) \leq p(\tilde{x}^{(j)})$ for negative monotonicity in k).

Non-monotonic decision-making can be justified, as many relationships between features and unobserved variables are themselves non-monotonic. For example, increasing income will usually improve creditworthiness, but exceptionally high incomes may correlate with volatile self-employment or irregular cash flow, which can legitimately lower repayment probability.

To check whether algorithmic decision-making conflicts with monotonicity rules, individuals need access to local feature-attribution information that identifies which features the model used and in which direction they influenced the prediction locally. One possible type of explanation to approximate such local influence directions is LIME. A LIME explanation estimates the linear contribution of each feature to the model’s decision in the neighborhood of the individual.

Definition 9 (LIME Explanation). Let $x \in \mathcal{X}$, then LIME collects the function values of the model $\hat{p}$ on points sampled from a local distribution π_x around x . A linear regression is then performed on the collected samples and function values, resulting in a linear model $g(x) = \sum_{i=1}^p w_i \cdot x_i$ approximating $\hat{p}$ in the local environment π_x . The local faithfulness of g to $\hat{p}$ on π_x is given by

$$\text{LF}(g, \hat{p}) := \mathbb{E}_{\pi_x}[(g(\pi_x) - \hat{p}(\pi_x))^2].$$

The following theorem shows that LIME explanations, given LIME coefficients w_k point in the opposite direction of monotonicity rules, only provide sufficient information to establish somewhere inaccuracy. However, as the counterexample shows, LIME explanations are insufficient to render decisions (somewhere) epistemically contestable.

Theorem 2. Let π be the local environment⁶ around $\tilde{x}^{(i)}$, with $\pi = (\pi_1, \dots, \pi_r)$ jointly independent random variables, and let π be constrained by $\mathbb{P}_\pi(|\pi - \tilde{x}^{(i)}| > \epsilon) = 0$. Assume that the monotonicity intuition applies positively (respectively negatively) to feature k and that the LIME explainer yields $w_k < 0$ (respectively $w_k > 0$). If the local faithfulness of g satisfies $\text{LF}(g, \hat{p}) < \mathbb{V}_{\pi_k}[w_k \pi_k]$, then the model is somewhere inaccurate.

⁶For notational simplicity, we abbreviate the local environment here as π rather than $\pi_{\tilde{x}^{(i)}}$.

Other, more direct, explanation methods for detecting monotonicity conflicts include individual conditional expectation curves (Goldstein et al. 2015) or prediction gradients (Wang et al. 2024), both of which also allow for the detection of somewhere inaccuracies. We chose LIME here to illustrate monotonicity conflicts because it is most commonly used in contestability-relevant domains such as lending or hiring.

4.3 Detecting Reason Conflicts With Anchors

Reason⁷ rules take the form: *If a person has certain feature values on a subset of relevant features, this provides sufficient grounds for making the decision.* Humans often justify their decisions by providing *reasons*, aspects of the specific case that motivated the decision (Raz 1975). These reasons can then be challenged by questioning their validity or by highlighting other reasons that override them.

Definition 10 (Reason rule). If an individual i has feature values satisfying conditions A , then this provides sufficient grounds to make the decision $d(\tilde{x}^{(i)}) = 1$ rather than $d(\tilde{x}^{(i)}) = 0$, or vice versa.

Some reason rules adopted by humans are in fact incorrect. For example, while many people assume that a very high salary suffices to justify a positive credit decision, substantial debt or an unstable job history may outweigh a high salary, thereby justifying a negative decision.

To assess whether algorithmic decision-making conflicts with reason rules, individuals need access to local decision rules that identify which minimal sets of feature values the model treats as sufficient grounds for its prediction. We believe that Anchors are the best available approach to approximate such locally sufficient conditions. An anchor identifies a set of feature constraints that, when satisfied, render the model’s prediction locally invariant with high probability in the surrounding neighborhood.

Definition 11 (Anchor Explanation). Let $x \in \mathcal{X}$, an *anchor* is a rule B (a set of feature conditions) such that, for points sampled from a local distribution π_x that satisfy B , the decision remains stable with high *precision*. The precision of an anchor B is defined as $\text{prec}(B) := \mathbb{E}_{\pi_x|B}[\mathbf{1}_{\hat{d}(x)=\hat{d}(\pi_x)}] = 1 - \delta$ for some $\delta \in [0, 1]$.

The following theorem shows that Anchor explanations, given the Anchor indicates the opposite decision as the rea-

⁷This reflects a specific conception of reasons that ties them to sufficiency for particular decisions. An alternative view treats reasons as factors that push a decision in a given direction (Snedegar 2018). On this account, good reasons are factors that, when present, increase the probability of the unobserved variable of interest toward the realized decision. However, this view faces a key difficulty: what counts as a good reason is often highly context-dependent and sensitive to specific counterfactual feature values, making it hard to formulate as a general rule. Moreover, even if suitable local rules could be specified, the most promising auditing method, counterfactual Shapley values (Albini et al. 2022), would still be inadequate, as it relies heavily on model predictions for data points far from the instance being explained and thus cannot even reliably detect somewhere inaccuracies.

son rule and both have sufficient overlap, can provide sufficient information to establish somewhere contestability. However, as the counterexample shows, they do not render decisions epistemically contestable.

Theorem 3. *Let i be an individual to whom the reason rule A applies implying a positive (respectively negative) decision and the anchor B applies, leading to a negative (respectively positive) decision with precision $\text{prec}(B) = 1 - \delta$ and $\mathbb{P}_{\pi_x}(B) > 0$, and let π be constrained by $\mathbb{P}_{\pi}(|\pi - \tilde{x}^{(i)}| > \epsilon) = 0$. If $\mathbb{P}_{\pi_x}(A \cap B) > \delta$, then the decision is somewhere contestable.*

5 Evidence that Allows for Normative and Strong Epistemic Contestability

We have shown that conflicts between human intuitions and standard XAI explanations reveal errors somewhere; however, as argued, such errors are insufficient to justify overturning a decision. What kind of evidence, then, would enable individuals to assess and demonstrate that a decision is normatively or epistemically contestable? In the following, we examine three types of evidence that can be sufficient to warrant reversal: predictive multiplicity, incorrect feature values, and ignored overruling evidence.

5.1 Predictive Multiplicity Reflects the Chances for Epistemic Contestability

Machine learning is, at its core, a tool of inductive inference. Given a dataset, e.g., containing creditworthy and non-creditworthy applicants, the goal is to learn a model that is able to distinguish loan applicants who will repay their loan from those who will not. Like any other tool of inductive inference, machine learning requires inductive assumptions to generalize to novel instances.

Modelers incorporate these inductive assumptions into the modeling pipeline through design choices that reflect their domain knowledge, always with the goal of best estimating the true underlying relationship p . Model classes, regularization techniques, and hyperparameters are selected based on this knowledge. In addition to such *intentional choices*, model selection is also influenced by *conventional choices*, such as the selection of a learning rate, or *arbitrary choices*, such as the selection of random seeds (Ganesh, Taik, and Farnadi 2025; Cavus, Woźnica, and Biecek 2025). Even though inductive assumptions play a central role in guiding the model selection process, the iron rule in machine learning is that only models with strong performance on a holdout set are ultimately selected (Hardt 2025).

As with all tools of inductive inference, this gives rise to a problem of *contrastive underdetermination* (Laudan 1990; Stanford 2023): multiple jointly incompatible models align equally well with both the observed data and the relevant domain knowledge, making it difficult to choose among them. In machine learning, this phenomenon has been called *model multiplicity* or the *Rashomon effect* (Breiman 2001; Black, Raghavan, and Barocas 2022). There always exists a multitude of prediction models that cannot be distinguished from an epistemic perspective: they result from the same model selection process and exhibit (almost) identical predictive

performance.⁸ The set of such models has been called the *Rashomon set*.

Definition 12 (Rashomon model and set). Let S denote the model selection process, and let $H_{D,\delta}$ denote a random variable describing the intentional, conventional, and arbitrary inductive choices, which are informed by the dataset D and the performance satisfaction level δ . For a given model $\hat{p} := S(H_{D,\delta}, D)$, we call

- $\hat{q}$ a *Rashomon model* relative to $\hat{p}$ if both result from the same model selection process applied to the same dataset and performance satisfaction level but differ in their inductive choices $H_{D,\delta}$; formally, $\hat{q} = S(H'_{D,\delta}, D)$ for some $H'_{D,\delta} \neq H_{D,\delta}$ with positive density; and
- $RS_{\hat{p}} := \{\hat{q}_1, \dots, \hat{q}_n\}$ a *Rashomon set* if it consists of n randomly drawn Rashomon models, formally induced by i.i.d. sampled inductive choices $H_{D,\delta}^{(i)} \sim H_{D,\delta}$ with $\hat{q}_i = S(H_{D,\delta}^{(i)}, D)$.

Since all Rashomon models surpass the performance satisfaction level, they exhibit similar performance. Indeed, similar performance is often used as a defining criterion for Rashomon models (Fisher, Rudin, and Dominici 2019; Watson-Daniels, Parkes, and Ustun 2023; Ganesh, Taik, and Farnadi 2025). We argue, however, that arising from the same model-selection process provides a more appropriate definition, as it ensures not only comparable compatibility with the data across Rashomon models, but also consistency with researchers' background knowledge as encoded in the inductive choices.

From a purely epistemic perspective, all Rashomon models must be regarded as on a par: they have equal claim to representing the ground truth dependence p . Despite this, they can differ substantially in their predictions for individual instances, a phenomenon known as *predictive multiplicity* (Marx, Calmon, and Ustun 2020).⁹ Since these models have equal epistemic justification, so do their predictions. If some predictions are positive while others are negative, this provides a meaningful basis for contesting the decision, as argued by Black, Raghavan, and Barocas (2022) and formalized in the following theorem.

The theorem shows that, unless the decision maker has good reason to regard the selected model as superior, predictive multiplicity yields a lower bound on the probability that a decision is epistemically contestable. In this way, predictive multiplicity offers a principled route to establishing epistemic contestability.

Theorem 4. *Let $\hat{p} := S(H_{D,\delta}, D)$ denote the prediction model used in decision $\hat{a}$, obtained from a random draw from $H_{D,\delta}$, and let $RS_{\hat{p}}$ denote a Rashomon set. Assume there is no measurement error in the feature values, i.e. $\tilde{x}^{(i)} = x^{(i)}$.*

⁸Indeed, non-epistemic values (Longino 2025), such as fairness or interpretability concerns, can guide model selection in such situations, an idea now widely discussed in the model multiplicity literature (Black, Raghavan, and Barocas 2022; Xin et al. 2022; Rudin et al. 2024; Meyer et al. 2025; Höltingen and Oliver 2025).

⁹A variety of global measures for predictive multiplicity have been proposed (Watson-Daniels, Parkes, and Ustun 2023).

Then the probability that $\hat{p}$ is epistemically incorrect for an individual i can be lower bounded by

$$\mathbb{P}_{H_{D,\delta}}(d(x^{(i)}) \neq \hat{d}(x^{(i)})) \geq \min\{\mathbb{P}_{H_{D,\delta}}(\hat{d}(x^{(i)})=1), \mathbb{P}_{H_{D,\delta}}(\hat{d}(x^{(i)})=0)\}.$$

We denote this lower bound by θ , which can be estimated via the minimum of the proportions of positive and negative decisions in the Rashomon set, i.e.

$$\hat{\theta}_{RS_{\hat{p}}} := \min\left\{\frac{1}{|RS_{\hat{p}}|} \sum_{m \in RS_{\hat{p}}} \mathbf{1}_{\hat{d}_m(x^{(i)})=1}, \frac{1}{|RS_{\hat{p}}|} \sum_{m \in RS_{\hat{p}}} \mathbf{1}_{\hat{d}_m(x^{(i)})=0}\right\}.$$

The estimator is asymptotically consistent, although it systematically underestimates θ in finite samples.¹⁰

Predictive multiplicity, and thus the prospects for successful epistemic contestability, is high in two scenarios: (1) when all predictions are close to τ but some lie above and others below, corresponding to high aleatoric uncertainty; and (2) when predictions are widely spread, with some far above and others far below τ , corresponding to high epistemic uncertainty (Hüllermeier and Waegeman 2021).

Policy implication. Individuals should be informed of the predictive multiplicity underlying their decision by reporting predictive multiplicity in their case. The respective Rashomon set should be randomly sampled from the space of plausible, high-performing models and constructed independently of the individual's feature values.

5.2 Wrong Feature Values Can Enable Epistemic Contestability

Algorithmic decision-making crucially relies on the measured feature values of individuals $\tilde{x}^{(i)}$. In the context of lending, some of these features are provided to the credit institution by the applicant herself, such as age, marital status, or the requested credit amount. Others stem from information the bank has access to in its databases, such as the individual's credit or transaction history. Finally, there is information that the bank actively gathers to assess the applicant's trustworthiness, e.g., public records and court judgments (Dansana et al. 2024).

During the processes of entering, transferring, or collecting data, errors can occur: the applicant's employment status may have changed; the credit history may contain redundant entries; or the collected public records may pertain to a different individual with the same name.¹¹ Such errors are often not visible to the individual, as she typically only knows the features she has entered herself, and even those may contain unnoticed mistakes.

To make reliable decisions, a solid epistemic basis is indispensable. A decision can only be supported through evidence if that evidence is correct (Conlee and Feldman 2004).

¹⁰The finite-sample downward bias can be bounded; see Clark (1961).

¹¹Similarly, feature values can be subject to uncertainty. For example, estimated monthly expenses or salary may vary substantially for self-employed or commission-based individuals.

Algorithmic decision-making thus inherits a standard familiar from theories of epistemic justification: if the informational premises on which an inference is based are defective, then the resulting belief or decision is epistemically unwarranted, even if the inferential tool itself is flawless (Wright 2003).

If the measured feature values $\tilde{x}^{(i)}$ differ substantially from the true values $x^{(i)}$, this discrepancy can lead to corresponding errors in predictions and decisions (Tal 2020). To assess whether this is the case, individuals should be granted *what-if access* to the model; that is, they should be allowed to test the model on a limited set of relevant input values.

Definition 13 (What-if). What-if provides individuals access to $\{\hat{p}(x) \mid x \in \mathcal{X}_{rel}\}$ for a limited set of relevant inputs $\mathcal{X}_{rel}$ specified by the individual subject to the decision.

The following theorem shows that incorrect feature values, when corrected values would reverse the decision, enable individuals to epistemically contest the decision. While the premise of epistemic correctness may appear strong, denying it would require decision makers to question the validity of the very model they rely on.

Theorem 5. Assume the decision model $\hat{d}$ is epistemically correct, i.e. $\hat{d} = d$. If individuals can prove that their true feature values are $x^{(i)}$ rather than $\tilde{x}^{(i)}$, and if $\hat{d}(x^{(i)}) \neq \hat{d}(\tilde{x}^{(i)})$, then the decision is epistemically contestable.

Policy implication. Individuals should be informed about their personal information used in the decision-making process. Additionally, they should be able to query the model using alternative inputs.

5.3 Overruling Evidence is a Means Towards Normative Contestability

Algorithmic decision-making reduces the epistemic basis of decisions to a fixed set of features to infer the unobserved property of interest, e.g., whether the loan will be repaid. While these features are carefully chosen, normatively incorrect decisions may still be made, even if the resulting decision closely approximates the epistemically correct one—even for optimal classifiers (Hastie et al. 2009).

Such a setting is highly static. By contrast, when humans make decisions, they flexibly incorporate additional information about other relevant feature values when these are presented to them. For example, if an individual is accused of cheating by her professor, she can provide further information to support her case, such as contemporaneous drafts, handwritten notes, or a version history demonstrating that the text was produced independently; this information is then taken into account in the professor's decision. In such cases, a decision may remain epistemically correct relative to the feature set X , but it is no longer epistemically correct once the additional features are considered. We can transfer the idea underlying such *overruling evidence* to the algorithmic case.

Definition 14 (Overruling evidence). We call features E *overruling evidence* for individual i if $d_X(x^{(i)}) \neq$

$d_{X,E}(x^{(i)}, e^{(i)})$, where $d_X := \mathbf{1}_{\mathbb{P}(Y=1|X) > \tau}$ and $d_{X,E} := \mathbf{1}_{\mathbb{P}(Y=1|X,E) > \tau}$.

Contesting a decision through overruling evidence differs fundamentally from other forms of contestation. Contestation based on XAI explanations, predictive multiplicity, or incorrect feature values all target the correctness of estimations or measurements within a fixed feature set X and aim to establish epistemic contestability relative to X . Overruling evidence, by contrast, accepts that the decision is epistemically correct relative to X but epistemically incorrect, and thus contestable, relative to the expanded feature set X, E . This invokes an epistemic meta-principle, namely the *Principle of Total Evidence* (Carnap 1947), according to which beliefs should always be based on all available evidence. This principle has strong normative support: it has been shown that new evidence can only increase, and never decrease, expected utility when the cost of information is negligible (Taroni et al. 2024; Schurz 2024).

An analogous result is shown below for our setting, demonstrating that fewer or equal normatively incorrect decisions are made when additional evidence is taken into account. This implies that decisions based on a smaller evidential basis are generally less justified than decisions based on all available evidence, thereby supporting the idea that overruling evidence should overturn a decision.

Theorem 6. *Assume that X and E describe two sets of features. Moreover, assume that the prediction models used are Bayes-optimal relative to X and to X, E , respectively, i.e. $\hat{p}_X = p_X = \mathbb{P}(Y | X)$ and $\hat{p}_{X,E} = p_{X,E} = \mathbb{P}(Y | X, E)$. Consequently, $\hat{d}_X = d_X$ and $\hat{d}_{X,E} = d_{X,E}$. Then, the expected number of normatively contestable decisions relative to X is greater than the expected number of normatively contestable decisions relative to X, E , i.e.*

$$\mathbb{E}_{Y,X,E}[d^* \neq \hat{d}_{X,E}] \leq \mathbb{E}_{Y,X}[d^* \neq \hat{d}_X].$$

Contesting a decision by means of overruling evidence is often the closest we can get to normative contestability, since the property of interest Y is typically unobservable. For example, it is not possible at decision time to observe whether the loan will ultimately be repaid. However, even when overruling evidence sometimes exists, this is not necessarily always the case; consider a loan applicant for whom all observable evidence speaks against repayment, but who would later win the lottery and repay the loan. Rendering such a normatively incorrect decision epistemically contestable would require Laplace’s demon (Frigg et al. 2014).

Overruling evidence is appealing in theory but difficult to implement in practice. Decision makers typically do not have access to an approximation of $d_{X,E}$ for arbitrary evidence E . This leads to a reversal of roles (Brennan-Marquez and Henderson 2019): individuals must explain why evidence E changes their repayment probability sufficiently to warrant their desired decision. In practice, this often results in a negotiation process between the decision maker and the individual, in which both parties attempt to reach an agreement (Lipsky 2010). Particularly when the deployed decision model exhibits low predictive performance¹² or when

uncertainty for a given individual is high, decision makers should have a strong incentive to take overruling evidence into account.

Policy implication. Individuals should be allowed to submit additional evidence concerning features that were not considered in the original decision-making process. This is particularly critical when the deployed predictive model has low predictive performance or when the decision involves high uncertainty.

6 Individuals and the Right to Contestation in the GDPR

So far, we have taken a purely *normative perspective*, focusing on the question of which decisions *should* be overturned from an epistemic standpoint. We argued that, to enable individuals to contest decisions, they should gain access to certain types of evidence, namely: information about predictive multiplicity, the feature values used in their case, API access to query the model with a limited number of alternative inputs, and the ability to provide additional feature information. We now shift from this normative approach to a *legal perspective*. Focusing on European legislation, we ask whether individuals have legal rights to access such evidence in the context of algorithmic contestability. A comprehensive answer to this question would require engaging with the intricacies of multiple EU legal instruments and is therefore beyond the scope of this paper. Nevertheless, we provide an initial discussion of how our proposed sources of evidence could relate to the GDPR, in particular Article 22 GDPR, one of the most prominent legal provisions governing individuals’ rights to contest automated decision-making.

Contestability forms the “backbone” of Article 22 GDPR (Bayamlioğlu 2022). It is laid down in Article 22(3) GDPR and applies where a decision that produces legal effects concerning the data subject or similarly significantly affects them is based solely on automated processing. Such decisions are prohibited in principle. They are permitted only where one of the exceptions in Article 22(2) GDPR applies, and even then suitable safeguards must be in place under certain circumstances. A recent Court of Justice of the European Union (CJEU) judgment (C-634/21, 7 December 2023, ECLI:EU:C:2023:957) ruled that, in the context of credit scoring, the concept of a “decision” within the meaning of Article 22 GDPR is to be understood broadly. However, the judgment did not address whether a person can “test the system” (Stewart 2026).

Interestingly, despite the centrality of the right to contest within Article 22, it remains “barely articulated” in the GDPR (Kaminski and Urban 2021). While the right to contest a decision is specifically laid down in Article 22(3) GDPR, Article 22(1) GDPR could, in a broader sense, also be interpreted as providing a right to object to being subject to automated decision-making, thereby extending the concept beyond strict contestation. So far, Article 15(1)(h) (together with Art. 22 GDPR) has primarily been interpreted by

¹²Low predictive performance can indicate that relevant features

XAI researchers as promoting a right to explanations which in turn should empower data subjects to contest decisions (Selbst and Powles 2018) (see also C-203/22 Dun & Bradstreet, 27 February 2025, ECLI:EU:C:2025:117 on explainability requirements in the GDPR).

For us it is an open question whether Article 22 GDPR can and should be interpreted more broadly. Some scholars argue that Article 22 GDPR could be reinterpreted to grant a right to verify systems and probe decisions (Stewart 2026). We agree with these directions. First, the GDPR specifies also that individuals have the right to express their point of view, aligning with our recommendation that individuals should be able to provide additional feature information to support their case. Second, the right to an explanation is not limited to current XAI solutions but encompasses any type of evidence that allows individuals to better understand the decision-making process (Vredenburg 2022). This may include what-if access through an API for alternative input values. Similarly, this should arguably include the uncertainty in the decision for their particular case, which can be expressed through predictive multiplicity.

One might argue that in contexts like loan approvals, providing the types of information we recommend could impose significant demands on decision makers. Access to predictive multiplicity would require training multiple models with equal performance, all of which must meet stringent fairness and accountability standards. Providing individuals with their complete data and API access may expose proprietary company information (Sokol and Flach 2019). Moreover, responding to all additional feature information submitted by individuals could be costly and burdensome (see also Article 12(5) GDPR and Article 15(4) GDPR).

However, without access to these types of information, individuals lack the necessary means to exercise their right to contestation, raising the possibility that the right effectively becomes meaningless. In contexts where decision makers cannot provide sufficient support for contestation, it may be questionable whether the decision should be automated at all. Importantly, XAI explanations alone, as we have shown, are insufficient to meet the demands of the right to contest, and it remains questionable whether XAI explanations alone should be considered an adequate effort by decision makers to fulfill the duties arising from Article 22 GDPR.

7 Discussion

In this paper, we conceptualized algorithmic contestability as the natural counterpart to traditional algorithmic recourse. Instead of recommending actions that change a decision, contestability concerns the provision of evidence that allows individuals to overturn a decision. We showed that the evidence provided by common XAI methods, together with human intuitions, allows only for the detection of errors somewhere and is insufficient to warrant decision reversal. Consequently, we moved beyond XAI and examined other sources of evidence that are, from a normative perspective, sufficient to overturn algorithmic decisions. We focused on predictive multiplicity, wrong feature values, and overruling evidence. For each, we grounded our analysis in epistemological theory and demonstrated that these forms of evidence indeed

allow decisions to be overturned. We concluded by arguing that current legislation in the European Union, in particular the GDPR, already grants individuals access to these kinds of evidence that enable them to overturn decisions.

Ethical Reasons for Contesting. Our work focused on epistemic grounds for contestation; however, there can also be *ethical reasons* for contestation. For instance, because the data may reflect societal biases or discrimination, using epistemically correct decisions as a standard becomes problematic, as these decisions are based on observational data influenced by such biases (Green 2022; Barocas, Hardt, and Narayanan 2023). Similarly, a model may rely on protected attributes in its decision-making that are predictive of the target but nevertheless discriminate against certain groups. Excluding such attributes is often insufficient to address this issue due to the existence of *proxy* variables (Dwork et al. 2012). Since these ethical concerns are already widely discussed in the fairness community, we deliberately focused on the thus-far overlooked epistemic aspects of contestability. In practice, decisions can of course be contested on both ethical and epistemic grounds.

Contesting When Individuals and Decision Makers are Adversarial. For future work, we believe that incorporating a stakeholders' perspective would be highly fruitful. In our setup, individuals can contest decisions by appealing to the ethical norms to which the decision maker commits. Thus, in our account, the goals of the involved parties are aligned: individuals who are likely to repay their loans aim to convince the decision maker that they will do so, while individuals who did not cheat aim to convince the decision maker that they acted honestly. Such situations are ideal, as contestability then consistently leads to better decisions (Cohen et al. 2023). Many real-world settings, however, are adversarial, and the interests of individuals and decision makers are not aligned (Bordt et al. 2022). In such contexts, many open questions remain: Who provides the evidence for contestability, and who mediates between the individual and the decision maker?

The Limits of XAI Explanations for Contestability. Our analysis was confined to a subset of human intuitions and XAI explanations. Clearly, there are further intuitive rules that humans adhere to; consider, for example, *causal-relevance rules*, according to which only features that are causally relevant to the unobserved property should affect the decision (Sloman and Lagnado 2015; Ye et al. 2024). Likewise, other XAI explanations, such as Shapley interaction values (Muschalik et al. 2024), global feature effect or importance measures (Molnar 2020), or approaches from mechanistic interpretability (Sharkey et al. 2025), could be analyzed with respect to their value for contestability. We believe that even if the analysis were extended along these lines, our core claim would remain: XAI explanations alone are insufficient to enable individuals to overturn decisions and must therefore be complemented by other forms of evidence.

Acknowledgments

We thank Tom Sterkenburg and the participants of the MCMP colloquium for valuable feedback on the

manuscript.

This research was supported by the German Research Foundation (DFG), project number 511917847; by the Cluster of Excellence “Machine Learning — New Perspectives for Science” (EXC 2064/1, project number 390727645); and by the Carl Zeiss Foundation through the CZS Center for AI and Law.

Generative AI Disclosure Statement

All authors confirm that generative AI tools were used only for minor support in manuscript preparation. Specifically, we used ChatGPT (version 5), Gemini (version 3) and Claude Opus 5 to assist with grammar and style editing, as well as with LaTeX drafting of boxes and tables. No text, claims, or references were generated by these tools. All authors take full responsibility for the originality, accuracy, and integrity of the manuscript. For coding, Claude Code (Opus 4.6) was used.

References

- Adadi, A.; and Berrada, M. 2018. Peeking inside the black-box: a survey on explainable artificial intelligence (XAI). *IEEE access*, 6: 52138–52160.
- Albini, E.; Long, J.; Dervovic, D.; and Magazzeni, D. 2022. Counterfactual shapley additive explanations. In *Proceedings of the 2022 ACM conference on fairness, accountability, and transparency*, 1054–1070.
- Aler Tubella, A.; Theodorou, A.; Dignum, V.; and Michael, L. 2020. Contestable black boxes. In *International Joint Conference on Rules and Reasoning*, 159–167. Springer.
- Alfrink, K.; Keller, I.; Kortuem, G.; and Doorn, N. 2023. Contestable AI by design: Towards a framework. *Minds and Machines*, 33(4): 613–639.
- Anders, B. A. 2023. Is using ChatGPT cheating, plagiarism, both, neither, or forward thinking? *Patterns*, 4(3).
- Barocas, S.; Hardt, M.; and Narayanan, A. 2023. *Fairness and machine learning: Limitations and opportunities*. MIT Press.
- Bayamlioğlu, E. 2022. The right to contest automated decisions under the General Data Protection Regulation: Beyond the so-called “right to explanation”. *Regulation & Governance*, 16(4): 1058–1078.
- Birnbaum, M. H. 2019. Violations of monotonicity in judgment and decision making. In *Choice, decision, and measurement*, 73–100. Routledge.
- Black, E.; Raghavan, M.; and Barocas, S. 2022. Model multiplicity: Opportunities, concerns, and solutions. In *Proceedings of the 2022 ACM conference on fairness, accountability, and transparency*, 850–863.
- Bordt, S.; Finck, M.; Raidl, E.; and Von Luxburg, U. 2022. Post-hoc explanations fail to achieve their purpose in adversarial contexts. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 891–905.
- Breiman, L. 2001. Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Statistical science*, 16(3): 199–231.
- Brennan-Marquez, K.; and Henderson, S. E. 2019. Artificial intelligence and role-reversible judgment. *J. Crim. L. & Criminology*, 109: 137.
- Cabitza, F.; Campagner, A.; and Basile, V. 2023. Toward a perspectivist turn in ground truthing for predictive computing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 6860–6868.
- Carnap, R. 1947. On the application of inductive logic. *Philosophy and phenomenological research*, 8(1): 133–148.
- Cavus, M.; Woźnica, K.; and Biecek, P. 2025. The Role of Hyperparameters in Predictive Multiplicity. arXiv:2503.13506.
- Clark, C. E. 1961. The greatest of a finite set of random variables. *Operations Research*, 9(2): 145–162.
- Cohen, I. G.; Babic, B.; Gerke, S.; Xia, Q.; Evgeniou, T.; and Wertenbroch, K. 2023. How AI can learn from the law: putting humans in the loop only on appeal. *npj Digital Medicine*, 6(1): 160.
- Conee, E.; and Feldman, R. 2004. *Evidentialism: Essays in epistemology*. Clarendon Press.
- Dandl, S.; Molnar, C.; Binder, M.; and Bischl, B. 2020. Multi-objective counterfactual explanations. In *International conference on parallel problem solving from nature*, 448–469. Springer.
- Dansana, D.; Patro, S. G. K.; Mishra, B. K.; Prasad, V.; Razak, A.; and Wodajo, A. W. 2024. Analyzing the impact of loan features on bank loan prediction using Random Forest algorithm. *Engineering Reports*, 6(2): e12707.
- Department of Industry, Innovation and Science. 2019. Artificial Intelligence: Australia’s Ethics Framework.
- Douglas, H. 2000. Inductive risk and values in science. *Philosophy of science*, 67(4): 559–579.
- Douglas, H. E. 2009. *Science, policy, and the value-free ideal*. University of Pittsburgh Press.
- Dwork, C.; Hardt, M.; Pitassi, T.; Reingold, O.; and Zemel, R. 2012. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, 214–226.
- Fisher, A.; Rudin, C.; and Dominici, F. 2019. All models are wrong, but many are useful: Learning a variable’s importance by studying an entire class of prediction models simultaneously. *Journal of Machine Learning Research*, 20(177): 1–81.
- Freiesleben, T. 2022. The intriguing relation between counterfactual explanations and adversarial examples. *Minds and Machines*, 32(1): 77–109.
- Frigg, R.; Bradley, S.; Du, H.; and Smith, L. A. 2014. Laplace’s demon and the adventures of his apprentices. *Philosophy of Science*, 81(1): 31–59.
- Ganesh, P.; Taik, A.; and Farnadi, G. 2025. Systemizing Multiplicity: The Curious Case of Arbitrariness in Machine Learning. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 8, 1032–1048.
- Gigerenzer, G. 2023. *The intelligence of intuition*. Cambridge University Press.

- Gigerenzer, G.; and Gaissmaier, W. 2011. Heuristic decision making. *Annual review of psychology*, 62(2011): 451–482.
- Goldstein, A.; Kapelner, A.; Bleich, J.; and Pitkin, E. 2015. Peeking inside the black box: Visualizing statistical learning with plots of individual conditional expectation. *Journal of Computational and Graphical Statistics*, 24(1): 44–65.
- Green, B. 2022. Escaping the impossibility of fairness: From formal to substantive algorithmic fairness. *Philosophy & Technology*, 35(4): 90.
- Hardt, M. 2025. The Emerging Science of Machine Learning Benchmarks. Online at <https://mlbenchmarks.org>. Manuscript.
- Hastie, T.; Tibshirani, R.; Friedman, J.; et al. 2009. *The elements of statistical learning*. Springer.
- Henin, C.; and Le Métayer, D. 2021. A framework to contest and justify algorithmic decisions. *AI and Ethics*, 1(4): 463–476.
- Hirsch, T.; Merced, K.; Narayanan, S.; Imel, Z. E.; and Atkins, D. C. 2017. Designing contestability: Interaction design, machine learning, and mental health. In *Proceedings of the 2017 Conference on Designing Interactive Systems*, 95–99.
- Höltgen, B.; and Oliver, N. 2025. Reconsidering Fairness Through Unawareness From the Perspective of Model Multiplicity. In *Proceedings of the 5th ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, 241–255.
- Hüllermeier, E.; and Waegeman, W. 2021. Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine learning*, 110(3): 457–506.
- Kaminski, M. E.; and Urban, J. M. 2021. The right to contest AI. *Columbia Law Review*, 121(7): 1957–2048.
- Kang, E. B. 2023. Ground truth tracings (GTT): On the epistemic limits of machine learning. *Big Data & Society*, 10(1): 20539517221146122.
- Karimi, A.-H.; Barthe, G.; Schölkopf, B.; and Valera, I. 2022. A survey of algorithmic recourse: contrastive explanations and consequential recommendations. *ACM Computing Surveys*, 55(5): 1–29.
- Kolodner, J. 2014. *Case-based reasoning*. Morgan Kaufmann.
- König, G.; Fokkema, H.; Freiesleben, T.; Mandler-Dünner, C.; and von Luxburg, U. 2025. Performative validity of recourse explanations. *Advances in Neural Information Processing Systems*, 39.
- König, G.; Freiesleben, T.; and Grosse-Wentrup, M. 2023. Improvement-focused causal recourse (ICR). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 11847–11855.
- Kusner, M. J.; Loftus, J.; Russell, C.; and Silva, R. 2017. Counterfactual fairness. *Advances in neural information processing systems*, 30.
- Laudan, L. 1990. Demystifying underdetermination. *Minnesota studies in the philosophy of science*, 14(1990): 267–297.
- Lipsky, M. 2010. *Street-level bureaucracy: Dilemmas of the individual in public service*. Russell Sage Foundation.
- Lipton, Z. C. 2018. The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue*, 16(3): 31–57.
- Longino, H. 2025. The Social Dimensions of Scientific Knowledge. In Zalta, E. N.; and Nodelman, U., eds., *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Spring 2025 edition.
- Lyons, H.; Velloso, E.; and Miller, T. 2021. Conceptualising contestability: Perspectives on contesting algorithmic decisions. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1): 1–25.
- Mansi, G.; Karusala, N.; and Riedl, M. 2025. Legally-Informed Explainable AI. *arXiv preprint arXiv:2504.10708*.
- Marx, C.; Calmon, F.; and Ustun, B. 2020. Predictive multiplicity in classification. In *International conference on machine learning*, 6765–6774. PMLR.
- Maxwell, W.; and Dumas, B. 2023. Meaningful XAI based on user-centric design methodology: Combining legal and human-computer interaction (HCI) approaches to achieve meaningful algorithmic explainability. Available at SSRN 4520754.
- Meyer, A. P.; Kim, Y.-S.; D’Antoni, L.; and Albarghouthi, A. 2025. Perceptions of the Fairness Impacts of Multiplicity in Machine Learning. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 1–15.
- Molnar, C. 2020. *Interpretable machine learning*. Lulu.com.
- Moreira, C.; Palatkina, A.; Braca, D.; Walsh, D. M.; Leihn, P. J.; Chen, F.; and Hubig, N. C. 2025. Explainable AI Systems Must Be Contestable: Here’s How to Make It Happen. *arXiv preprint arXiv:2506.01662*.
- Mulligan, D. K.; Kluttz, D.; and Kohli, N. 2019. Shaping our tools: Contestability as a means to promote responsible algorithmic decision making in the professions. Available at SSRN 3311894.
- Muschalik, M.; Baniecki, H.; Fumagalli, F.; Kolpaczki, P.; Hammer, B.; and Hüllermeier, E. 2024. shapiq: Shapley interactions for machine learning. *Advances in Neural Information Processing Systems*, 37: 130324–130357.
- Pawelczyk, M.; Lakkaraju, H.; and Neel, S. 2023. On the privacy risks of algorithmic recourse. In *International Conference on Artificial Intelligence and Statistics*, 9680–9696. PMLR.
- Raz, J. 1975. Reasons for action, decisions and norms. *Mind*, 84(336): 481–499.
- 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.
- Ribeiro, M. T.; Singh, S.; and Guestrin, C. 2018. Anchors: High-precision model-agnostic explanations. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.

- Rudin, C.; Zhong, C.; Semenova, L.; Seltzer, M.; Parr, R.; Liu, J.; Katta, S.; Donnelly, J.; Chen, H.; and Boner, Z. 2024. Position: amazing things come from having many good models. In *Proceedings of the 41st International Conference on Machine Learning*, 42783–42795.
- Sarra, C. 2020. Put dialectics into the machine: protection against automatic-decision-making through a deeper understanding of contestability by design. *Global Jurist*, 20(3): 20200003.
- Schmude, T.; Yurrita, M.; Alfrink, K.; Goff, T. L.; and Viard, T. 2025. "Two Means to an End Goal": Connecting Explainability and Contestability in the Regulation of Public Sector AI. *arXiv preprint arXiv:2504.18236*.
- Scholbeck, C. A.; Molnar, C.; Heumann, C.; Bischl, B.; and Casalicchio, G. 2019. Sampling, intervention, prediction, aggregation: a generalized framework for model-agnostic interpretations. In *Joint European conference on machine learning and knowledge discovery in databases*, 205–216. Springer.
- Schurz, G. 2024. The Principle of Total Evidence: Justification and Political Significance. *Acta Analytica*, 39(4): 677–692.
- Selbst, A.; and Powles, J. 2018. "Meaningful information" and the right to explanation. In *conference on fairness, accountability and transparency*, 48–48. PMLR.
- Selbst, A. D.; and Barocas, S. 2018. The intuitive appeal of explainable machines. *Fordham L. Rev.*, 87: 1085.
- Sharkey, L.; Chughtai, B.; Batson, J.; Lindsey, J.; Wu, J.; Bushnaq, L.; Goldowsky-Dill, N.; Heimersheim, S.; Ortega, A.; Bloom, J.; et al. 2025. Open problems in mechanistic interpretability. *arXiv preprint arXiv:2501.16496*.
- Sloman, S. A.; and Lagnado, D. 2015. Causality in thought. *Annual review of psychology*, 66(1): 223–247.
- Snedegar, J. 2018. Reasons for and reasons against. *Philosophical Studies*, 175(3): 725–743.
- Sokol, K.; and Flach, P. 2019. Counterfactual explanations of machine learning predictions: opportunities and challenges for AI safety. In *2019 AAAI Workshop on Artificial Intelligence Safety, SafeAI 2019*. CEUR Workshop Proceedings.
- Stanford, K. 2023. Underdetermination of Scientific Theory. In Zalta, E. N.; and Nodelman, U., eds., *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Summer 2023 edition.
- Stewart, M. 2026. Beyond Explanation: Evidentiary Rights for Algorithmic Accountability. In *The 2026 ACM Conference on Fairness, Accountability, and Transparency*, 1066–1097.
- Tal, E. 2020. Measurement in Science. In Zalta, E. N., ed., *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2020 edition.
- Taroni, F.; Aitken, C.; Bozza, S.; and Juchli, P. 2024. The principle of total evidence revisited. *Law, Probability and Risk*, 23(1): mgae011.
- Tversky, A.; and Kahneman, D. 1974. Judgment under Uncertainty: Heuristics and Biases: Biases in judgments reveal some heuristics of thinking under uncertainty. *science*, 185(4157): 1124–1131.
- Upadhyay, S.; Joshi, S.; and Lakkaraju, H. 2021. Towards robust and reliable algorithmic recourse. *Advances in Neural Information Processing Systems*, 34: 16926–16937.
- Ustun, B.; Spangher, A.; and Liu, Y. 2019. Actionable recourse in linear classification. In *Proceedings of the conference on fairness, accountability, and transparency*, 10–19.
- Vaccaro, K.; Karahalios, K.; Mulligan, D. K.; Kluttz, D.; and Hirsch, T. 2019. Contestability in algorithmic systems. In *Companion Publication of the 2019 Conference on Computer Supported Cooperative Work and Social Computing*, 523–527.
- Venkatasubramanian, S.; and Alfano, M. 2020. The philosophical basis of algorithmic recourse. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, 284–293.
- Von Kügelgen, J.; Karimi, A.-H.; Bhatt, U.; Valera, I.; Weller, A.; and Schölkopf, B. 2022. On the fairness of causal algorithmic recourse. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, 9584–9594.
- Vredenburg, K. 2022. The right to explanation. *Journal of Political Philosophy*, 30(2): 209–229.
- Wachter, S.; Mittelstadt, B.; and Russell, C. 2017. Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *Harv. JL & Tech.*, 31: 841.
- Wang, Y.; Zhang, T.; Guo, X.; and Shen, Z. 2024. Gradient based feature attribution in explainable AI: A technical review. *arXiv preprint arXiv:2403.10415*.
- Watson-Daniels, J.; Parkes, D. C.; and Ustun, B. 2023. Predictive multiplicity in probabilistic classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 10306–10314.
- Wright, C. 2003. Some reflections on the acquisition of warrant by inference. *New essays on semantic externalism and self-knowledge*, 57–77.
- Xin, R.; Zhong, C.; Chen, Z.; Takagi, T.; Seltzer, M.; and Rudin, C. 2022. Exploring the whole rashomon set of sparse decision trees. *Advances in neural information processing systems*, 35: 14071–14084.
- Ye, W.; Zheng, G.; Cao, X.; Ma, Y.; and Zhang, A. 2024. Spurious correlations in machine learning: A survey. *arXiv preprint arXiv:2402.12715*.
- Yurrita, M.; Draws, T.; Balayn, A.; Murray-Rust, D.; Tintarev, N.; and Bozzon, A. 2023. Disentangling fairness perceptions in algorithmic decision-making: the effects of explanations, human oversight, and contestability. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–21.
- Yurrita, M.; Verma, H.; Balayn, A.; Alfrink, K.; Gadiraju, U.; and Bozzon, A. 2025. Identifying Algorithmic Decision Subjects' Needs for Meaningful Contestability. *Proceedings of the ACM on Human-Computer Interaction*, 9(7): 1–29.