

Beyond Technocratic XAI: The Who, What & How in Explanation Design

Ruchira Dhar¹, Stephanie Brandl², Ninell Oldenburg³, Anders Søgaard^{1,3}

¹ Department of Computer Science, University of Copenhagen

² Center for Social Data Science, University of Copenhagen

³ Department of Philosophy, University of Copenhagen

rudh@di.ku.dk, stephanie.brandl@sodas.ku.dk, niol@hum.ku.dk, soegaard@di.ku.dk

Abstract

The field of Explainable AI (XAI) offers a wide range of techniques for making complex models interpretable. Yet, in practice, generating meaningful explanations is a context-dependent task that requires intentional design choices to ensure accessibility and transparency. This paper reframes explanation as a situated design process—an approach particularly relevant for practitioners involved in building and deploying explainable systems. Drawing on prior research and principles from design thinking, we propose a three-part framework for explanation design in XAI: asking *Who* needs the explanation, *What* they need explained, and *How* that explanation should be delivered. We also emphasize the need for ethical considerations, including risks of epistemic inequality, reinforcing social inequities, and obscuring accountability and governance. By treating explanation as a sociotechnical design process, this framework encourages a context-aware approach to XAI that supports effective communication and the development of ethically responsible explanations.

Introduction

The field of Explainable AI (XAI) offers a wide array of methods to interpret and communicate the decisions made by complex models. Most works, however, focus on developing novel methods in the space (Ribeiro, Singh, and Guestrin 2016; Lundberg and Lee 2017), creating taxonomies to make users aware of different methods (Guidotti et al. 2018; Arrieta et al. 2020), and creating scientific frameworks for comparing and evaluating methods (Doshi-Velez and Kim 2017; Vilone and Longo 2021). However, when deployed in practice, producing an explanation is a context-sensitive process that requires careful thought and design (Huuskonen and Korteniemi 1992; Kim et al. 2023; Sokol and Vogt 2024; Wolf 2019; Wolf and Ringland 2020). In domains like healthcare, finance, or criminal justice, explanations are not just tools for understanding but mechanisms for transparency, accountability, and trust-building. When misaligned with user needs or social context, explanations can mislead and reinforce existing inequities (Gerlings, Jensen, and Shollo 2021; Rane, Choudhary, and Rane 2023; Chaudhary 2024). This leads us to an important question: *how*

could someone best approach the task of generating an explanation in context? While existing research highlights the need for human-centric XAI development (Edmonds et al. 2019; Ehsan et al. 2022, 2024b) and design awareness for bridging sociotechnical gaps (Ehsan et al. 2023, 2024a), there remains a gap in suggesting comprehensive end-to-end usable frameworks that can be easily applied to the design of an XAI process across deployment contexts.

What is needed is not just a toolbox of techniques, but a structured way to think about the explanation process itself, from intention to implementation. For XAI practitioners working to translate complex model behavior into accessible and trustworthy explanations, this shift is especially critical (Das and Rad 2020; Asghari et al. 2021; Munn and Pitman 2022). To illustrate this, let us consider a scenario in which an XAI practitioner is tasked with explaining how a machine learning system detects early signs of a disease. Before designing the explanation, they must ask:

- **WHO** is going to use the explanation? A doctor or a patient? Their clinical background will determine what kind of explanation is appropriate.
- **WHAT** needs to be explained? Should the explanation give features that influenced the model’s decision? Or the false positive rate of all predictions in general? The depth and scope of the content depend on the user’s role and decision-making needs.
- **HOW** should the explanation be delivered? Should it be accompanied by a brief textual summary or a confidence score? What kind of language should explanations use? The delivery format must align with the user’s workflow and cognitive preferences.

Moreover, once the explanation is delivered, an important question arises: does it support responsible and fair decision-making? Scholars have increasingly emphasized that explanation is not ethically neutral—it involves decisions about who gets to understand a system, what kind of knowledge is privileged, and how accountability is distributed (McDermid et al. 2021). In high-stakes domains like healthcare or finance, ethical adequacy means clarity on one hand but also enabling accountability and trust on the other (Gerlings, Jensen, and Shollo 2021). Such considerations can have a significant impact on how an explanation is designed for users within different contexts.

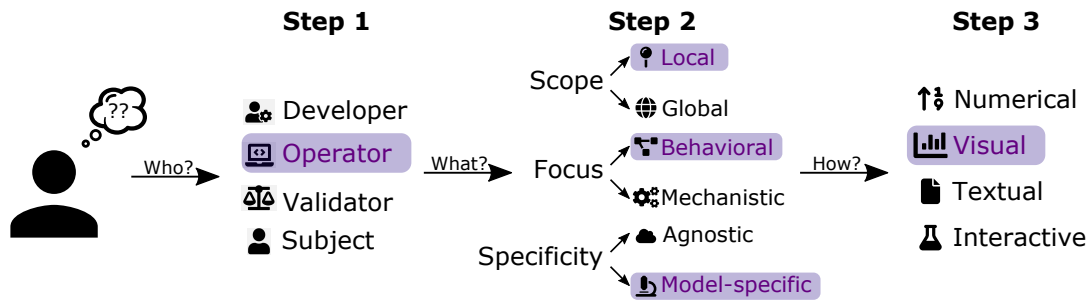


Figure 1: An Overview of our proposed XAI Design Process

This simplistic example captures a key point: explanations are always for someone, about something, and must be designed accordingly. Taking inspiration from design thinking (Brown et al. 2008; Brown 2015) and research advocating human-centric explanations (Wang et al. 2019; Ehsan and Riedl 2020; Ehsan et al. 2021; Frank and George 2024; Oldenburg and Sogaard 2025), we propose a framework for treating explanation as a sociotechnical practice. This also complements recent work in XAI evaluation that emphasizes the importance of user-centered metrics and human factors (Liao, Gruen, and Miller 2020; Ehsan et al. 2024b; Frank and George 2024). While our methodological discussion primarily focuses on widely used XAI approaches in natural language processing (NLP), the framework we present is adaptable to a broad range of AI systems and application contexts.

In *Background & Related Work*, we discuss some essential terminology used in the paper. We also discuss related work in the field that has served as inspiration and is relevant to the premise of our work. In *Design Process in XAI*, we address each step of the process individually and provide key consideration points in each step. In *Ethics in XAI*, we highlight some key considerations for reflecting on ethical aspects of the decisions made in the design process. In *Reframing XAI as Design Process*, we summarize the process so far, and in *Design Process in Action*, we also take on a concrete example for illustrating the process with an imagined use-case. We give a step-wise instantiation of our design process, and in each step, provide ethical reflections on design choices made during the process.

By shifting from method-centric thinking to process-oriented design, we aim to support practitioners in building explanation systems that are not only technically robust but also contextually appropriate and ethically sound.

Background & Related Work

Terminology Explainability is commonly understood as the ability to articulate why a model produced a given output in a way that is accessible to human users, often emphasizing post hoc techniques that support understanding of model outputs (Arrieta et al. 2020; Wiegrefe et al. 2022; Dwivedi et al. 2023; Kunz and Kuhlmann 2024). Interpretability, by contrast, is sometimes defined more narrowly as the degree to which a model’s internal mechanics can be directly understood, typically by experts and often through inherently

transparent model architectures or diagnostic tools (Doshi-Velez and Kim 2017). More recently, Mechanistic Interpretability (MI) has emerged as an umbrella term for another category of interpretability methods that are mostly aimed at interpreting model outputs in terms of their internal representations and architectures (Bereska and Gavves 2024). The “traditional” XAI community and the more recent MI community so far operate mostly in parallel. Saphra and Wiegrefe (2024) made a first step towards disentangling the two communities. While the distinction of explainability and interpretability has guided parts of the literature (Miller 2019), both terms are often used interchangeably in the literature.

In this paper, we do not maintain a strict separation between explainability and interpretability. Instead, we treat explanation as a concept that spans the entire pipeline from interrogating black-box models to communicating justifications that are meaningful to stakeholders. This includes both technical introspection and communicative delivery.

XAI Methods and Their Limits A wide range of technical methods for XAI have been proposed, including feature attribution (e.g., Integrated Gradients, Layer-wise Relevance Propagation), surrogate models (e.g., LIME, SHAP), and example-based explanations (e.g., counterfactuals). These methods aim to increase transparency but often fall short in practice due to their reliance on technical assumptions that do not align with users’ mental models or real-world constraints (Alvarez-Melis and Jaakkola 2018). Moreover, they are designed without considering the interpretive needs or capacities of *specific* users. Recent work in mechanistic interpretability (Saphra and Wiegrefe 2024) offers a complementary view focused on opening the “black box” by analyzing internal model structure, but largely remains developer-oriented. In this paper, we will try to close this gap by including the most prominent explainability approaches into our design framework.

Human-Centered & Design-Oriented XAI Recognizing the limits of purely technical explanations, researchers have explored user-centered and participatory approaches to XAI (Abdul et al. 2018; Kaur et al. 2020; Ehsan et al. 2022; Ehsan and Riedl 2020; Ehsan et al. 2023, 2024b,a). These include user studies on explanation efficacy (Jacovi et al. 2023), comparing model explanations to either human attention in the form of eye-tracking (Brandl et al. 2024) or





Stakeholder Type	Demographic	Domain Knowledge	Technical Knowledge	Goal
 Developers	Researchers Engineers	Low	High	Analyze, debug, and improve model behavior
 Operators	Doctors Bankers	High	Low	Use model outputs to support decisions
 Validators	Auditors Compliance Officers	Medium	Medium	Assess fairness, robustness, and regulatory compliance
 Subjects	Patients Customers	Low	Low	Understand, contest, or build trust in decisions that affect them

Table 1: Answering the *WHO*: Categorizing stakeholders to support explanation design in XAI

human annotations (Jakobsen, Cabello, and S ogaard 2023), and interactive explanation interfaces (Wexler et al. 2020; Hohman et al. 2019). While there is some work on design-oriented XAI (Wolf 2019; Wolf and Ringland 2020), there is a lack of general systematic frameworks to guide explanation design across contexts. Our work builds on this tradition, proposing a principled design framework grounded in user-centric thinking.

Ethical Perspectives on Explanation Explanation is not a value-neutral activity but involves decisions about whose understanding matters, what kinds of knowledge are privileged, and how accountability is distributed (Binns 2018; Selbst et al. 2019). Scholars in Science Technology Studies and Human-Computer Interaction (HCI) have highlighted the risks of epistemic injustice (Langton 2010), design bias (Costanza-Chock 2020), and abstraction traps (Selbst et al. 2019) in XAI systems. Our work positions explanation as a sociotechnical process that must grapple with these ethical and normative tensions, offering a framework for designing explanations that are not only useful but ethical.

Design Process in XAI

In this section, we introduce our main contribution: a human-centered framework for designing explanations in XAI, organized around three key questions: *Who, What, and How* as outlined in Figure 1. This framework encourages a shift from focusing solely on algorithmic methods to considering the broader context in which explanations are created and used. It is also the basis for delineating the ethical accounts of explanation design in the next section.

Answering the WHO

The first question in explanation design is simple: *Who is the explanation for?* Unlike most technical work in XAI that assumes a generic user, human-centered design demands that we attempt to identify and differentiate among the stakeholders engaged with or affected by an AI system (Miller 2019; Decker et al. 2023).

Existing work has tried to come up with criteria for identifying stakeholders in the explainability space—ranging from chronological order in a traditional machine learning pipeline (Dwivedi et al. 2023), the need for explanations (Calderon and Reichart 2025), and legal considerations

(Mansi, Karusala, and Riedl 2025). For applied contexts, however, there are few requirements to be kept in mind:

- **Reliance on professional titles** can create false dichotomies where multiple stakeholder groups might have similar goals (e.g., developers, theorists, and data scientists can all have conflated goals of understanding the model better in general).
- **Stakeholder classification** needs to be based on goals but also be fine-grained enough to identify and differentiate individual in real world scenarios. Calderon and Reichart (2025) clubs together significantly different stakeholders like enterprises and end-users (who can have different interests in XAI methods) while creating a theoretically vague class of “society” which potentially includes all stakeholders at large.

We take inspiration from Habibullah et al. (2024) which proposes a user taxonomy based on a rigorous literature review of the field. In addition, we provide factors of consideration such as demographics, domain knowledge, technical knowledge, and goals that can help distinguish actors in real-world scenarios. Table 1 summarizes four groups: Developers, Validators, Operators, and Subjects. This framing draws on both life cycle roles and interpretive purpose, while highlighting differences in access and accountability.

- **Developers** require explanations to understand, debug, and refine model internals. Explanations must prioritize technical faithfulness and granularity.
- **Operators** use model outputs to support real-time decision-making in high-stakes settings. Explanations must be context-sensitive, actionable, and efficient.
- **Validators** such as auditors or compliance officers seek transparency to assess fairness, risk, or adherence to norms. This group can also include advocacy groups who have similar goals of ensuring fairness. Explanations here must be comprehensible, structured, and aligned with regulatory standards.
- **Subjects** are individuals impacted by model decisions. They may not interact directly with the model, but their rights and outcomes are shaped by it. Explanations here must be accessible to empower recourse.

It is important to note that the roles are not always mutually exclusive, but that a single individual may occupy

multiple positions across the AI lifecycle depending on real-world contexts of deployment. However, our framework is meant to provide guidance in such situations instead of enforcing definitive and rigid dichotomies. By articulating the knowledge constraints and distinct goals of different stakeholders, this classification supports the design of explanation methods that are meaningfully aligned with their intended users. For instance, a developer may take on the role of validator during model audits. Despite these overlaps, each role reflects a distinct mindset shaped by specific constraints, expectations, and forms of accountability. Designing effective explanations requires attending to these mindsets as meaningful user perspectives, not just organizational labels (Dhanorkar et al. 2021; Ehsan et al. 2024b).

Answering the WHAT

Once the relevant stakeholder has been identified, the next critical design decision is: *What should be explained?* This question shapes not only the type of explanation to be provided but also its relevance, utility, and impact. In practice, it is not enough to generate technically sound explanations. Effective XAI must align with stakeholder goals, cognitive capacity, and operational context. An explanation that targets the wrong aspect of model behavior risks misleading its audience or becoming practically unusable.

To support better-aligned design decisions, we classify existing XAI methods along four axes that reflect their real-world implications (Madsen, Reddy, and Chandar 2022; Dwivedi et al. 2023; Calderon and Reichart 2025): (1) **explanation scope**, (2) **explanation focus**, (3) **model specificity**, and (4) **operational cost**. We see these dimensions as design-relevant levers that help match methods to the needs of different stakeholders, introduced earlier. Our discussion centers primarily on post hoc explanation methods typically applied to textual AI systems like LLMs (Wang et al. 2024), i.e., models that are typically black-box in nature and do not offer intrinsic or self-explaining capabilities (Loyola-Gonzalez 2019). Furthermore, these four design axes are broadly applicable and can be extended to guide explanation strategies across a wider range of AI models and domains, even though our emphasis lies on methods relevant to contemporary NLP systems.

Explanation Scope This axis refers to the level of generality at which a method operates. **Local methods** provide explanations for individual predictions. These are crucial in high-stakes applications where specific decisions must be justified, such as loan approvals or medical diagnoses. Popular local techniques include gradient-based methods like integrated gradients (Sundararajan, Taly, and Yan 2017), some surrogate methods like LIME (Ribeiro, Singh, and Guestrin 2016), or counterfactual methods (Dai et al. 2022).

In contrast, **Global methods** aim to explain the overall behavior of a model across many inputs. They are well-suited for understanding patterns the model has learned and are often useful in development or validation settings. Examples include some surrogate methods like SHAP (Lundberg and Lee 2017) and global surrogate modelling (Apley and Zhu 2020; Zhu et al. 2022), probing (Wexler et al. 2020),

circuit tracing (Olah et al. 2020; Conmy et al. 2023), and sparse autoencoders (Sharkey, Braun, and Millidge 2022). Some other popular methods also include partial dependence plots (Friedman 2001; Greenwell 2017), individual conditional expectation (Goldstein et al. 2015), etc.

From a stakeholder lens, global methods are typically valuable for *developers* and *validators* who are focused on the model’s general behavior, while local methods better serve *operators* and *subjects* who are concerned with case-specific outcomes.

Explanation Focus Different explanation methods shed light on different aspects of a model. Are we illuminating which parts of the input were most influential on the model’s output, or are we uncovering the internal mechanisms, such as features, representations, or circuits, that the model uses to compute its decision? **Behavioral methods**, often referred to as “post-hoc” methods, aim to explain how inputs influence outputs, often by observing model behavior under perturbations or through feature attribution (Retzlaff et al. 2024; Cesarini et al. 2024). These include gradient-based feature attribution methods like Integrated Gradients (Sundararajan, Taly, and Yan 2017), but also Surrogate Models like SHAP (Lundberg and Lee 2017) and LIME (Ribeiro, Singh, and Guestrin 2016), counterfactuals, and probing techniques.

By contrast, **mechanistic methods** attempt to uncover the internal structures, representations, or circuits that drive model behavior. This emerging paradigm, often associated with mechanistic interpretability (Olah et al. 2020; Nanda et al. 2023), includes tools such as neuron-level attribution (Geiger et al. 2021), sparse autoencoders (Elhage et al. 2022), and circuit tracing techniques (Conmy et al. 2023).

The choice of explanation focus should be guided by what the stakeholder needs to understand to achieve their goal. *Developers*, for example, often benefit from mechanistic explanations that reveal how internal components such as neurons, layers, or circuits contribute to model behavior, enabling them to debug or refine the system. *Validators*, such as auditors or fairness reviewers, may require both behavioral and mechanistic insights to assess robustness, compliance, or potential bias. For *operators* and *subjects*, however, behavioral explanations tend to be more intuitive and actionable, particularly example-based or counterfactual methods. They offer clear links between inputs, outputs, and possible alternatives. Crucially, this design choice does not just affect what the explanation reveals, but it has important consequences for how the explanation is interpreted and used in sensitive decision-making scenarios (Jesus et al. 2021; Vale, El-Sharif, and Ali 2022).

Model Specificity Explanation methods also differ significantly in how much internal access they require to the model being explained and how specific any explanations are to the model at hand. This axis, often referred to as model specificity (Darias, Díaz-Agudo, and Recio-Garcia 2021; Sandu and Trausan-Matu 2022; Letrache and Ramdani 2023), is a key aspect in determining the relevance of explanations.

The **model-specific** approaches rely on internals like gradients (Sundararajan, Taly, and Yan 2017), representations

Method	Scope	Focus	Specificity	Best-fit Stakeholders
Grad.-based Feat. Attrib.	Local	Behavioral	Specific	Developers, Validators
Surrogate Methods	Local/Global	Behavioral	Agnostic	Developers, Validators
Counterfactuals	Local	Behavioral	Agnostic	Operators, Validators, Subj.
Probing	Global	Behavioral	Specific	Validators
Circuit Tracing	Global	Mechanistic	Specific	Developers
Sparse Autoencoders	Global	Mechanistic	Specific	Developers

Table 2: Answering the *WHAT*: Taxonomy of explanation methods across design-relevant axes. We group methods here into categories. Gradient-based feature attribution includes methods like LRP, Integrated Gradients, Saliency Mapping, and Gradient \times Input. SHAP and LIME fall under the umbrella of (local) surrogate models.

for probing (Wexler et al. 2020), attention weights for circuits (Conmy et al. 2023), or neuron activations and features in sparse autoencoders (Massidda et al. 2023). These methods are tightly coupled to the model’s architecture and learning dynamics, allowing for high-resolution insights into how a particular model processes information. For example, techniques like Integrated Gradients or LRP compute feature attributions based on backpropagated gradients, while transformer-specific methods can trace attention flow or patch activations to localize functionality. Since model specific approaches can access the internal model internal components, the explanations are grounded in the model’s internal structure and behavior (?). However, they require full access to the model and its internals, which is not always feasible when the model is closed-source or where privacy concerns are paramount (Spartalis, Semertzidis, and Daras 2023; Ezzeddine 2024).

On the other hand, **model-agnostic** methods treat the model as a black box, relying only on inputs and outputs. These include methods relying on surrogate models like LIME and (kernel) SHAP or counterfactuals. While more widely applicable and well-suited for settings where the underlying model is proprietary, model-agnostic methods may offer less faithful approximations and struggle with complex or highly non-linear models (Molnar et al. 2020).

From a design perspective, *developers* and *validators* working in trusted environments may prefer model-specific methods, while *operators* and *subjects*, who often face opaque third-party models, depend on model-agnostic approaches. For them, model-agnostic methods offer essential transparency, even if less granular. Selecting an explanation method, therefore, must align not only with what the model can support, but also with who the user is and what kind of interaction is realistically possible.

Operational Cost Finally, practical deployment requires attention to the cost of explanation, both in terms of **time latency** and **compute load**. Perturbation-based methods like LIME (Ribeiro, Singh, and Guestrin 2016) and kernel SHAP (Lundberg and Lee 2017; Roshan and Zafar 2022) often involve hundreds of model queries, resulting in noticeable delays and high memory use. Gradient-based methods are lighter because they reuse model internals, but can become costly when explaining many inputs in batch or when dealing with massive models. Mechanistic methods are often

the most computationally intensive, involving large-scale activation patching, neuron tracing, or circuit discovery that may require dedicated infrastructure (Luccioni et al. 2025).

These costs are not just a matter of engineering but affect who can realistically benefit from XAI. In low-latency environments, e.g., mobile traffic classification (Nascita et al. 2021) or anomaly detection (Gummadi, Napier, and Abdallah 2024), lightweight behavioral methods may be the only viable option. In resource-constrained or sustainability-focused settings, compute-heavy techniques may be ethically or environmentally untenable (Stojkovic et al. 2024; Luccioni, Jernite, and Strubell 2024; Luccioni et al. 2025). Thus, cost-conscious design is essential, and some recent work has made progress towards attempts at including cost considerations of XAI in model design (Jean-Quartier et al. 2023).

Cost considerations are a critical yet often overlooked aspect of explanation design in XAI. In general, *developers* and *validators* may accept higher cost for deeper analysis, but *operators* and *subjects* typically need lightweight, fast, and low-resource explanations.

Taken together, these four axes provide a practical lens for selecting explanation methods that align with the diverse needs, capacities, and constraints of different stakeholders. Table 2 offers an illustrative mapping of commonly used XAI methods across three of the axes: scope, focus, and model specificity. We exclude operational cost from this table, as it is highly variable and depends on factors such as model architecture, input size, and the number of instances to be explained. Similarly, the stakeholder-method mappings should not be read as definitive recommendations; rather, they serve as indicative starting points. The aim of this illustration is not to give prescriptive matches, but starting points for informed design decisions.

Answering the HOW

The final design question becomes: *How should the explanation be delivered?* But this question is more than a technical decision about format. It is an act of epistemic mediation, i.e., deciding how knowledge is shaped, prioritized, and framed for specific ends. As Miller (2019) notes, explanations are not mere outputs but social performances: constructing narratives, eliciting trust or doubt, and reflecting assumptions about what kinds of understanding “count.”

Modality	Description	Example Tools / Techniques	Best-fit Stakeholders
↑ Numerical	Confidence scores, feature importances, statistical summaries	SHAP values, LIME values	👤 Developers
📊 Visual	Plots, saliency maps, heatmaps, dashboards	Grad-CAM Plots, LIME visualizations	👤 Dev., 🧑 Val., 👤 Op.
📄 Textual	Explanatory text, simplified justifications, summaries	Natural language rationales, Rules, Counterfactuals, Chain-of-Thought Explanations	🧑 Val., 👤 Op., 👤 Subj.
🔧 Interactive	Exploratory tools allowing user interaction with model behavior	What-if tools, Causal simulators	👤 Op., 👤 Subj.

Table 3: Answering the *HOW*: Typical Modalities of Explanation Delivery

In this view, explanation delivery is fundamentally a design problem. It involves selecting the right **format**, **modality**, and **structure** to support cognitive alignment between the model and its users. Recent work also indicates that explanation format significantly influences user trust, comprehension, and reliance (Kulesza et al. 2013; Liao, Gruen, and Miller 2020; Ehsan et al. 2024b). An explanation that is accurate but poorly delivered may be ignored, misunderstood, or even mistrusted. Effective explanation thus requires careful attention to how insights are surfaced.

Common explanation formats fall into four broad categories: **Numerical**, **Visual**, **Textual**, and **Interactive**. Each offers affordances as well as limitations, omissions, and implicit ideologies about the nature of understanding.

Numerical Explanations These explanations use quantitative indicators such as scores, weights, or ranks to convey feature importance or model confidence. Common methods include gradient-based feature attribution like Integrated Gradients or surrogate methods like LIME, SHAP, etc. Numerical formats are compact, expressive, and precise, especially useful for *developers* or *validators* who are comfortable with abstract metrics. However, they assume statistical fluency and often lack narrative support, which makes them rather inaccessible to non-experts. Moreover, they often mask uncertainty under a veneer of objectivity, privileging technical legibility over narrative intelligibility. In contexts like healthcare or public services, this can reinforce epistemic asymmetries between technical experts and impacted individuals (Wolf and Ringland 2020).

Visual Explanations Visual methods leverage graphical encodings like heatmaps, saliency maps, and attention plots to illustrate what the model “sees” or attends to. Examples include Grad-CAM (Selvaraju et al. 2020), or attention visualization in transformers (Belrose et al. 2023). These methods offer immediate, intuitive understanding and are especially powerful for *operators* and *subjects* who benefit from fast glanceable insight.

Note that visual explanations are often generated from the numerical explanations that are generated first. Thus, they can oversimplify or obscure causal mechanisms, and may lead users to overinterpret the output (Jain and Wallace 2019; Lyu, Apidianaki, and Callison-Burch 2024).

Textual Explanations Textual explanations transform model reasoning into natural language, either as post-hoc rationalizations or integrated reasoning chains. Variants include *explain-then-predict* (Lei, Barzilay, and Jaakkola 2016; Camburu et al. 2018), *predict-then-explain* (Marasovic et al. 2022), and *chain-of-thought* prompting in LLMs. These formats excel at accessibility and narrative coherence, particularly for *subjects* and non-technical stakeholders. However, natural language explanations may sacrifice faithfulness or omit uncertainty, creating a false sense of understanding if not critically designed (Agarwal, Tanneru, and Lakkaraju 2024; Edmonds et al. 2019).

As recent work shows, textual explanations are prone to hallucination, simplification, and selective emphasis (Agarwal, Tanneru, and Lakkaraju 2024). In general, language invites narrative closure and a premature sense of “knowing” that may conceal model fallibility or normative bias. When these explanations are deployed to justify high-stakes decisions (e.g., denial of loans), their rhetorical polish can dangerously undermine accountability rather than support it.

Interactive Explanations Interactive formats empower users to engage with the explanation, exploring what-if scenarios, adjusting inputs, or navigating between explanation layers. Tools like Google’s *What-If Tool* (Wexler et al. 2020) or *Gamut* (Hohman et al. 2019) blend visual, numerical, and textual elements into a responsive interface. Interactive systems offer flexible, layered access to explanation, letting users manipulate inputs, simulate counterfactuals, or explore uncertainty (Oldenburg and Søgaard 2025). These are especially attractive to *developers*, *validators*, and skilled *operators* navigating complex models.

It’s important to note that these formats are not mutually exclusive and can also be combined. A single explanation can be translated across modalities, e.g., SHAP values rendered as bar charts, tables, or verbal summaries. The challenge hereby lies in choosing the right format for the user’s context, skill level, and goals.

In human-centered XAI, explanation delivery is consequential (Ehsan et al. 2024b). When poorly rendered, the same logic may confuse or mislead; when well-crafted, it can inform, empower, and build trust (Kulesza et al. 2013). *Developers* often prefer low-level numerical or interactive formats that support debugging while *validators*

benefit from mixed-modal explanations that document decisions while revealing latent patterns (Naiseh et al. 2023). *Operators* such as clinicians or analysts rely on fast, interpretable visual or interactive formats to make quick decisions (Scharowski et al. 2023; Brankovic et al. 2025). *Subjects* require natural language or counterfactuals that connect model decisions to personal impact, often with minimal cognitive load (Ehsan et al. 2024b).

Finally, even within stakeholder groups, users vary. A senior developer may grasp low-level metrics, while a new team member may need higher-level structure. Moreover, users often suffer from the illusion of understanding, i.e., “good-enough” comprehension: believing they understand an explanation when their mental model is incomplete (Kaur et al. 2024; Jacovi et al. 2023). To prevent this problem, it is important to consider how explanation formats not only shape what is explained but also what is understood.

Ethics in XAI

Explanations in AI are often framed as a technical remedy for opacity, or as an add-on for transparency or trust. However, this framing obscures the ethical stakes. Decisions about the who, what, and how of explanations are not neutral, but shape users’ ability to know, challenge, and contest automated decisions (Stone and Mittelstadt 2025). In this section, we highlight three dimensions that demand ethical attention in the design of explainable AI systems: epistemic inequality, social inequality, and accountability.

Epistemic Inequality

Explanations are not distributed equally across stakeholders. Developers and validators typically gain access to rich model internals and tooling, while affected individuals such as patients, applicants, or defendants often receive vague rationales or simplified narratives (Crawford et al. 2016; Ananny and Crawford 2018). This selective legibility reinforces what Fricker calls *epistemic injustice* (Fricker 2017), in which marginalized groups are denied full participation in knowledge practices that impact them. Recent work in AI ethics and design has highlighted how information asymmetry in ML systems can deepen structural exclusion (Edwards and Veale 2017; Veale and Binns 2017; Veale and Brass 2019). Explanation frameworks must grapple with this asymmetry, not only by identifying stakeholder needs, but by resisting the normalization of partial or inaccessible explanations for those with the most at stake.

Social Inequality

Explanation practices are shaped by and can reinforce broader social hierarchies. Explanations often function as curated narratives that reflect institutional priorities more than user empowerment, where users are frequently given just enough information to comply with a decision, but not enough to challenge it (Eubanks 2018). Without critical reflection, explainability becomes a tool of legitimation, not justice. Designers must be attentive to how explanation choices reflect and reproduce systemic inequality in sociotechnical settings (Selbst et al. 2019; Barocas, Selbst, and Raghavan 2020; Smart and Kasirzadeh 2024).

Accountability and Governance

Designing explanations is not just about making models interpretable but about making systems accountable (Doshi-Velez et al. 2017; Kim, Park, and Suh 2020; Novelli, Taddeo, and Floridi 2024). Explanations should help users understand model outputs *and* surface the values, assumptions, and trade-offs embedded in the broader sociotechnical system (Ehsan and Riedl 2020; Alpsancar et al. 2024).

Accountability has become a matter of governance rather than simply an ethical ideal (Mökander and Floridi 2022; Lechterman 2022; Atoum 2025). As AI systems increasingly mediate access to housing, credit, healthcare, and public services, governments and institutions are beginning to impose regulatory frameworks to ensure these systems remain transparent, fair, and contestable. Recent policy developments such as the EU AI Act (Veale and Zuiderveen Borgesius 2021), Canada’s AIDA (Artificial Intelligence and Data Act) (Muhammad and Yow 2023), and algorithmic auditing mandates in cities like New York (Mökander et al. 2022) reflect a growing consensus: explanations must not only serve users, but also serve governance functions by enabling oversight, documentation, and redress.

In this context, explanation is part of a broader infrastructure for institutional accountability beyond being a technical or cognitive aid. This includes designing explanations that can support audits (Zhang, Cho, and Vasarhelyi 2022), creating documentation that traces design decisions, surfacing model uncertainties, and anticipating failure modes. Effective explanation design must therefore support accountability at multiple levels: interpersonal (e.g., enabling users to ask “why?”), organizational (e.g., allowing teams to trace and justify outcomes), and systemic (e.g., enabling regulators and public to interrogate high-stakes decisions). Explanation should not only clarify what a model did, but also make visible how responsibility is distributed, how decisions can be contested, and how harms can be redressed.

Reframing XAI as Ethical Design Process

Calls to reframe XAI as a human-centered process are increasingly common, but often underspecified in practice (Wolf 2019; Wolf and Ringland 2020; Ehsan and Riedl 2020; Ehsan et al. 2021). We have argued that explanation in AI systems must not be seen as a technical afterthought or post hoc rationalization, but as a consequential act of communication embedded within sociotechnical systems. Having made clear the stakes of ethical accounts in XAI, in this section, we will detail how our presented framework can account for ethical considerations.

The **WHO** dimension centers stakeholder diversity, not as a checklist of user roles (developer, validator, operator, subject), but as a reminder that explanations must navigate conflicting priorities and asymmetrical access to interpretive resources. Who gets a convincing explanation is a question of epistemic justice. A system that delivers rich, layered explanations to developers but vague justifications to subjects risks reinforcing existing inequities (Lazar 2024).

The **WHAT** dimension highlights the contested nature of what is worth explaining. Whether one explains model

inputs, outputs, processes, or counterfactuals is not a neutral design choice. These decisions shape how responsibility is distributed, which forms of error are visible, and whose knowledge is considered legitimate. Moreover, the drive for “faithfulness” often privileges model-internal logic over user-centered relevance, obscuring the social or policy context in which decisions unfold.

The **HOW** dimension, as discussed in the previous section, involves the translation of technical reasoning into formats legible to human interpreters. But rather than being the end of the design process, the delivery is part of an ongoing negotiation between intelligibility, persuasion, and control. Explanation design is not simply about usability; it is about shaping trust, setting boundaries of critique, and deciding who is entitled to question the system.

Most importantly, reframing XAI as an ethical design process is a call for epistemic humility, social equity, and accountability. It requires recognizing that explanation is always situated, produced under constraints, shaped by norms, and open to strategic misuse. This perspective challenges practitioners to ask not just whether explanations are accurate, but whether they are just, aid understanding, and whose understanding they prioritize (Lazar 2024). And not just how they function, but how they might fail. In this light, XAI becomes less about “making models explainable” and more about making such systems accountable through design choices that are explicit, reflective, and open to critique.

Process Design in Action

To concretize the idea of explanation as process design, we walk through a structured example using the WHO, WHAT, and HOW framework. This thought experiment illustrates how explanation design unfolds under real-world constraints. Importantly, it also reveals how design decisions, though seemingly technical, are deeply ethical, reflecting choices about whose knowledge is privileged, what user assumptions are made, and which costs are deemed acceptable. Our example is meant as a guideline of implementing our framework and not a definitive reflection of exactly how the XAI process should look in deployment scenarios.

Case Study: Risk Prediction in Healthcare

Consider a predictive model in a hospital context. The model estimates the 30-day readmission risk for patients with chronic heart failure. It is based on features like age, lab results, comorbidities, medications, and prior hospitalizations.

The model already performs well on historical data. Now, it needs to support real-world decisions. How can an explanation model be constructed given the following points of interest?

- The intended user is a clinician.
- The explanation must support actionable interventions.
- The hospital has limited resources to invest in explanation tooling.

WHO: Identifying the Stakeholder (and Who Gets Excluded)

At first glance, the primary stakeholder is the clinician, an *Operator* making treatment decisions based on model outputs. This classification guides the design process:

- *Context*: A clinician receives a high-risk alert for a patient and must respond quickly.
- *Expertise*: Deep medical training, limited familiarity with ML.
- *Goal*: Understand the prediction well enough to decide whether to intervene.

Ethical Reflection: Centering the clinician in explanation design seems intuitive, given their decision-making role. However, this choice already shapes the epistemic boundaries of the system. By prioritizing the interpretive needs of the clinician, the design risks excluding the patient, the person most directly affected by the decision, from the explanation loop. This reflects a broader pattern of *epistemic inequality*, where those with institutional authority are granted access to richer, more actionable explanations, while subjects of the model are offered only minimal or no insight into how decisions are made (Fricker 2017). Furthermore, treating the clinician as a single, unified user archetype flattens important differences within the stakeholder group (Cabitza, Rasoini, and Gensini 2017). A senior cardiologist and an overburdened medical resident may differ significantly in cognitive capacity, trust in automation, and time available to engage with an explanation. Without attending to these intra-stakeholder variances, design may unintentionally reinforce existing hierarchies in clinical workflows, privileging some users over others (Horsky 2025). This step thus demands a reflection on whose interpretive needs are being centered, whose are sidelined, and what institutional power dynamics are being reproduced in the process.

WHAT: Defining What Gets Explained (and What Gets Omitted)

Designing the explanation content involves choosing from several dimensions:

- *Scope*: Local, patient-specific explanation is needed.
- *Focus*: Feature attribution—what factors drove this specific prediction.
- *Model access*: Full access permits use of model-specific techniques.
- *Constraints*: Only low-compute methods are viable.

This points us to a feature attribution method like SHAP (Lundberg and Lee 2017), which can decompose the prediction into additive contributions of input features. It may seem like SHAP values are not amenable to understanding of our relevant stakeholder here (since physicians may not understand how SHAP works) but this should be resolved in the next “HOW” step where the focus is on how we communicate the explanation in an accessible format.

Ethical Reflection: While SHAP satisfies the design criteria, it also encodes assumptions that often go unquestioned:

that causality can be approximated through additive feature contributions; that the model's internal logic is a sufficient proxy for human understanding; and that technical constraints (e.g., compute) take precedence over epistemic ones (e.g., explanation completeness). These trade-offs may be acceptable, but only if they are surfaced, not silently embedded in the tooling. Explanation choices can conceal how social biases are encoded in model predictions. Methods like SHAP focus on which features contributed to a specific output, but they often overlook whether those features reflect underlying disparities in the training data. For example, if prior hospitalizations or medication access correlate with socioeconomic status, the explanation may quietly reproduce systemic inequities without surfacing them. Without careful dataset curation and proactive bias assessments, explanations risk legitimizing discriminatory patterns as neutral technical outputs (Antoniadi et al. 2021; Saraswat et al. 2022; Albahri et al. 2023). Therefore, explanation design content must go beyond technical fidelity to actively interrogate how social hierarchies are embedded in both data and model logic. What is left unexplained matters for accountability. When explanations focus solely on local feature effects, they may satisfy transparency checkboxes but fall short of supporting audits, redress, or governance.

HOW: Communicating the Explanation (and Whose Cognitive Fit We Privilege)

Finally, we decide how to communicate these explanations effectively, tailoring the format to the stakeholder's needs and context. In this case, we have our clinician who operates in a high-stakes, time-sensitive environment and needs to interpret the explanation quickly and confidently.

- Clinicians are accustomed to structured health data (e.g., lab reports, vitals) and often rely on both data visualizations
- The SHAP method we chose is amenable to creating visualizations.

Given the clinician's workflow and cognitive constraints, a ranked bar chart of SHAP values may be appropriate:

Top drivers: Recent ER admission (+0.24), Low sodium (+0.18), High blood pressure (+0.15)

To enhance usability, we might supplement this with a natural language summary:

"This patient is predicted to be high-risk primarily due to a recent ER visit, low sodium levels, and multiple prior readmissions."

Ethical Reflection: While such explanations may appear intuitive, they often present the model's perspective as fact rather than construction. Moreover, designing for "cognitive fit" can reinforce existing power imbalances, optimizing for what the high-status stakeholder finds legible, rather than empowering broader participation in contesting the model's decision. We must also contend with the risk of false confidence: clinicians may over-trust these explanations without being aware of their limitations, such as missing confounders, spurious correlations, or non-causal reasoning paths. Ethically responsible communication must

therefore also include what the model cannot know, how it might fail, and where caution is warranted.

Further Notes

From an ethical perspective on accountability and governance, explanation should not be treated as a one-time output but as part of an ongoing system of oversight and responsibility (Smith-Renner et al. 2020). Once an explanation is delivered, it must be possible to trace how it was generated, how it influenced decisions, and what outcomes it contributed to. This requires concrete institutional practices such as:

- **Documentation:** To clearly state each explanation method's assumptions, limitations, and intended use cases (Bunn 2020).
- **Evaluation:** The XAI methods being used should always be evaluated and such evaluation must be available on request to all stakeholders for greater transparency (Mohseni, Zarei, and Ragan 2021; Agarwal et al. 2022; Nauta et al. 2023).
- **Feedback Mechanisms:** To allow clinicians to flag confusing or misleading explanations. Such feedback is also often useful for improving the XAI system itself (Förster et al. 2025).
- **Audit Trails:** To log when and how explanations are accessed and used in clinical decisions (Bagwe 2025).

Without such support, explanations risk offering performative transparency, appearing interpretable while failing to enable contestation, accountability, or institutional learning. Ethical XAI design must therefore embed explanations into broader workflows of institutions.

Conclusion and Future Work

This paper proposed a structured framework for explanation design in XAI, grounded in stakeholder-centered inquiry and ethical reflection. By centering the questions of *who* an explanation is for, *what* it should convey, and *how* it should be delivered, we aim to reframe XAI as a design challenge embedded in social context, not just a technical problem of model introspection. Rather than prescribing fixed rules, this framework offers a foundation for creating explanation systems that are more adaptive and attentive to real-world use. However, there still remain many directions for future work.

First, empirical validation must go beyond traditional metrics, e.g., satisfaction, trust scores (Alam and Mueller 2021; Brdnik, Podgorelec, and Šumak 2023) and engage with the lived experiences of stakeholders using XAI in practice (Hoffman et al. 2023). As recent work in HCI shows, explanation systems often produce unintended effects, such as misplaced trust, cognitive overload, or institutional drift from original design intent (Bhatt et al. 2020; Liao et al. 2022). Ethnographic methods (Wolf 2019) and participatory evaluation (Hashmati et al. 2024) will be essential to assess not only usability, but also which groups benefit from an explanation strategy, and which may be excluded, confused, or burdened by it.

Second, there is a pressing need for design tools and workflows that are both actionable and ethically robust. Toolkits for explanation should support inclusive participation during prototyping, as well as throughout deployment, and iteration (Ehsan et al. 2021). Such methods would ideally include checklists that flag risks of epistemic injustice (Fricker 2017), and interfaces that let users question or reject explanations (Sovrano and Vitali 2021). However, simplification can also conceal harm: participatory design must avoid becoming a token gesture, and future work should develop critical reflection practices that help teams recognize when efficiency or resource constraints are quietly narrowing whose needs are being served.

Third, explanation design must engage more directly with the material constraints and political realities that shape real-world AI systems. Resource limitations are not just technical boundaries; they reflect institutional priorities and power dynamics (Edwards and Veale 2017, 2018). As emerging regulation makes clear, technology will get increasingly tied to compliance and auditability (Mökander and Floridi 2022). Future research should explore how explanation quality is distributed across institutions (e.g., well-resourced vs underfunded hospitals) and stakeholders (e.g., clinicians vs patients), and investigate not only how to adapt explanations under constraint, but why those constraints exist, who imposes them, and whether they can or should be contested (Birhane et al. 2022a,b; Farrow 2023).

Ultimately, this framework invites the reader to rethink what explanation is and what it is for. Explanation shouldn't be treated as a static output or interface widget, but as an evolving negotiation among values, stakeholders, and institutions. Explanation should be understood not as a final product, but as part of a broader system of interaction, intervention, and responsibility. Future work should treat explanation as infrastructure that supports understanding, action, and accountability across real-world settings.

Acknowledgments

The authors would like to express their sincere gratitude to the anonymous reviewers for their valuable comments and constructive feedback. We also acknowledge the support of Pioneer Centre for AI under Grant No. P1, whose funding made this research possible.

References

Abdul, A.; Vermeulen, J.; Wang, D.; Lim, B. Y.; and Kankanhalli, M. 2018. Trends and trajectories for explainable, accountable and intelligible systems: An hci research agenda. In *Proceedings of the 2018 CHI conference on human factors in computing systems*, 1–18.

Agarwal, C.; Krishna, S.; Saxena, E.; Pawelczyk, M.; Johnson, N.; Puri, I.; Zitnik, M.; and Lakkaraju, H. 2022. Openxai: Towards a transparent evaluation of model explanations. *Advances in neural information processing systems*, 35: 15784–15799.

Agarwal, C.; Tanneru, S. H.; and Lakkaraju, H. 2024. Faithfulness vs. plausibility: On the (un) reliability of ex-

planations from large language models. *arXiv preprint arXiv:2402.04614*.

Alam, L.; and Mueller, S. 2021. Examining the effect of explanation on satisfaction and trust in AI diagnostic systems. *BMC medical informatics and decision making*, 21(1): 178.

Albahri, A. S.; Duham, A. M.; Fadhel, M. A.; Alnoor, A.; Baqer, N. S.; Alzubaidi, L.; Albahri, O. S.; Alamoodi, A. H.; Bai, J.; Salhi, A.; et al. 2023. A systematic review of trustworthy and explainable artificial intelligence in healthcare: Assessment of quality, bias risk, and data fusion. *Information Fusion*, 96: 156–191.

Alpsancar, S.; Buhl, H. M.; Matzner, T.; and Scharlau, I. 2024. Explanation needs and ethical demands: unpacking the instrumental value of XAI. *AI and Ethics*, 1–19.

Alvarez-Melis, D.; and Jaakkola, T. S. 2018. On the robustness of interpretability methods. *arXiv preprint arXiv:1806.08049*.

Ananny, M.; and Crawford, K. 2018. Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *new media & society*, 20(3): 973–989.

Antoniadi, A. M.; Du, Y.; Guendouz, Y.; Wei, L.; Mazo, C.; Becker, B. A.; and Mooney, C. 2021. Current challenges and future opportunities for XAI in machine learning-based clinical decision support systems: a systematic review. *Applied Sciences*, 11(11): 5088.

Apley, D. W.; and Zhu, J. 2020. Visualizing the effects of predictor variables in black box supervised learning models. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 82(4): 1059–1086.

Arrieta, A. B.; Díaz-Rodríguez, N.; Del Ser, J.; Bennetot, A.; Tabik, S.; Barbado, A.; García, S.; Gil-López, S.; Molina, D.; Benjamins, R.; et al. 2020. Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information fusion*, 58: 82–115.

Asghari, H.; Birner, N.; Burchardt, A.; Dicks, D.; Fassbender, J.; Feldhus, N.; Hewett, F.; Hofmann, V.; Kettemann, M. C.; Schulz, W.; et al. 2021. What to explain when explaining is difficult? An interdisciplinary primer on XAI and meaningful information in automated decision-making. *Alexander von Humboldt Institute for Internet and Society*.

Atoum, I. 2025. Revolutionizing AI Governance: Addressing Bias and Ensuring Accountability Through the Holistic AI Governance Framework. *International Journal of Advanced Computer Science & Applications*, 16(2).

Bagwe, C. 2025. Explainable AI (XAI) in Compliance Audits: Bridging the Gap Between AI and Regulatory Transparency. *International Journal of Scientific Engineering and Science*, 9: 137–144.

Barocas, S.; Selbst, A. D.; and Raghavan, M. 2020. The hidden assumptions behind counterfactual explanations and principal reasons. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, 80–89.

Belrose, N.; Furman, Z.; Smith, L.; Halawi, D.; Ostrovsky, I.; McKinney, L.; Biderman, S.; and Steinhardt, J. 2023.

- Eliciting latent predictions from transformers with the tuned lens. *arXiv preprint arXiv:2303.08112*.
- Bereska, L.; and Gavves, S. 2024. Mechanistic Interpretability for AI Safety - A Review. *Transactions on Machine Learning Research*. Survey Certification, Expert Certification.
- Bhatt, U.; Xiang, A.; Sharma, S.; Weller, A.; Taly, A.; Jia, Y.; Ghosh, J.; Puri, R.; Moura, J. M.; and Eckersley, P. 2020. Explainable machine learning in deployment. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, 648–657.
- Binns, R. 2018. Fairness in machine learning: Lessons from political philosophy. In *Conference on fairness, accountability and transparency*, 149–159. PMLR.
- Birhane, A.; Isaac, W.; Prabhakaran, V.; Diaz, M.; Elish, M. C.; Gabriel, I.; and Mohamed, S. 2022a. Power to the people? Opportunities and challenges for participatory AI. In *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, 1–8.
- Birhane, A.; Ruane, E.; Laurent, T.; S. Brown, M.; Flowers, J.; Ventresque, A.; and L. Dancy, C. 2022b. The forgotten margins of AI ethics. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 948–958.
- Brandl, S.; Eberle, O.; Ribeiro, T.; Søgaaard, A.; and Hollenstein, N. 2024. Evaluating Webcam-based Gaze Data as an Alternative for Human Rationale Annotations. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 6544–6556.
- Brankovic, A.; Cook, D.; Rahman, J.; Delaforce, A.; Li, J.; Magrabi, F.; Cabitza, F.; Coiera, E.; and Bradford, D. 2025. Clinician-informed XAI evaluation checklist with metrics (CLIX-M) for AI-powered clinical decision support systems. *npj Digital Medicine*, 8(1): 364.
- Brdnik, S.; Podgorelec, V.; and Šumak, B. 2023. Assessing perceived trust and satisfaction with multiple explanation techniques in XAI-enhanced learning analytics. *Electronics*, 12(12): 2594.
- Brown, T. 2015. When everyone is doing design thinking, is it still a competitive advantage. *Harvard Business Review*, 27.
- Brown, T.; et al. 2008. Design thinking. *Harvard business review*, 86(6): 84.
- Bunn, J. 2020. Working in contexts for which transparency is important: A recordkeeping view of explainable artificial intelligence (XAI). *Records Management Journal*, 30(2): 143–153.
- Cabitza, F.; Rasoini, R.; and Gensini, G. F. 2017. Unintended consequences of machine learning in medicine. *Jama*, 318(6): 517–518.
- Calderon, N.; and Reichart, R. 2025. On Behalf of the Stakeholders: Trends in NLP Model Interpretability in the Era of LLMs. In Chiruzzo, L.; Ritter, A.; and Wang, L., eds., *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 656–693. Albuquerque, New Mexico: Association for Computational Linguistics. ISBN 979-8-89176-189-6.
- Camburu, O.-M.; Rocktäschel, T.; Lukasiewicz, T.; and Blunsom, P. 2018. e-snli: Natural language inference with natural language explanations. *Advances in Neural Information Processing Systems*, 31.
- Cesarini, M.; Malandri, L.; Pallucchini, F.; Seveso, A.; and Xing, F. 2024. Explainable ai for text classification: Lessons from a comprehensive evaluation of post hoc methods. *Cognitive Computation*, 1–19.
- Chaudhary, G. 2024. Explainable Artificial Intelligence (xAI): Reflections on Judicial System. *Kutafin Law Review*, 10(4): 872–889.
- Conmy, A.; Mavor-Parker, A.; Lynch, A.; Heimersheim, S.; and Garriga-Alonso, A. 2023. Towards automated circuit discovery for mechanistic interpretability. *Advances in Neural Information Processing Systems*, 36: 16318–16352.
- Costanza-Chock, S. 2020. *Design justice: Community-led practices to build the worlds we need*. The MIT Press.
- Crawford, K.; Whittaker, M.; Elish, M. C.; Barocas, S.; Plasek, A.; and Ferryman, K. 2016. The AI now report. *The Social and Economic Implications of Artificial Intelligence Technologies in the Near-Term*, 2.
- Dai, X.; Keane, M. T.; Shaloo, L.; Ruelle, E.; and Byrne, R. M. 2022. Counterfactual explanations for prediction and diagnosis in XAI. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, 215–226.
- Darias, J. M.; Diaz-Agudo, B.; and Recio-Garcia, J. A. 2021. A Systematic Review on Model-agnostic XAI Libraries. In *ICCBR workshops*, 28–39.
- Das, A.; and Rad, P. 2020. Opportunities and challenges in explainable artificial intelligence (xai): A survey. *arXiv preprint arXiv:2006.11371*.
- Decker, T.; Gross, R.; Koebler, A.; Lebacher, M.; Schnitzer, R.; and Weber, S. H. 2023. The thousand faces of explainable AI along the machine learning life cycle: industrial reality and current state of research. In *International Conference on Human-Computer Interaction*, 184–208. Springer.
- Dhanorkar, S.; Wolf, C. T.; Qian, K.; Xu, A.; Popa, L.; and Li, Y. 2021. Who needs to know what, when?: Broadening the Explainable AI (XAI) Design Space by Looking at Explanations Across the AI Lifecycle. In *Proceedings of the 2021 ACM Designing Interactive Systems Conference*, 1591–1602.
- Doshi-Velez, F.; and Kim, B. 2017. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*.
- Doshi-Velez, F.; Kortz, M.; Budish, R.; Bavitz, C.; Gershman, S.; O'Brien, D.; Scott, K.; Schieber, S.; Waldo, J.; Weinberger, D.; et al. 2017. Accountability of AI under the law: The role of explanation. *arXiv preprint arXiv:1711.01134*.
- Dwivedi, R.; Dave, D.; Naik, H.; Singhal, S.; Omer, R.; Patel, P.; Qian, B.; Wen, Z.; Shah, T.; Morgan, G.; et al. 2023. Explainable AI (XAI): Core ideas, techniques, and solutions. *ACM Computing Surveys*, 55(9): 1–33.

- Edmonds, M.; Gao, F.; Liu, H.; Xie, X.; Qi, S.; Rothrock, B.; Zhu, Y.; Wu, Y. N.; Lu, H.; and Zhu, S.-C. 2019. A tale of two explanations: Enhancing human trust by explaining robot behavior. *Science Robotics*, 4(37): eaay4663.
- Edwards, L.; and Veale, M. 2017. Slave to the algorithm? Why a 'right to an explanation' is probably not the remedy you are looking for. *Duke L. & Tech. Rev.*, 16: 18.
- Edwards, L.; and Veale, M. 2018. Enslaving the algorithm: From a "right to an explanation" to a "right to better decisions"? *IEEE Security & Privacy*, 16(3): 46–54.
- 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.
- Ehsan, U.; Liao, Q. V.; Passi, S.; Riedl, M. O.; and Daumé III, H. 2024a. Seamless xai: Operationalizing seamless design in explainable ai. *Proceedings of the ACM on Human-Computer Interaction*, 8(CSCW1): 1–29.
- Ehsan, U.; Passi, S.; Liao, Q. V.; Chan, L.; Lee, I.-H.; Muller, M.; and Riedl, M. O. 2024b. The who in XAI: how AI background shapes perceptions of AI explanations. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 1–32.
- Ehsan, U.; and Riedl, M. O. 2020. Human-centered explainable ai: Towards a reflective sociotechnical approach. In *HCI International 2020-Late Breaking Papers: Multimodality and Intelligence: 22nd HCI International Conference, HCII 2020, Copenhagen, Denmark, July 19–24, 2020, Proceedings 22*, 449–466. Springer.
- Ehsan, U.; Saha, K.; De Choudhury, M.; and Riedl, M. O. 2023. Charting the sociotechnical gap in explainable AI: A framework to address the gap in XAI. *Proceedings of the ACM on human-computer interaction*, 7(CSCW1): 1–32.
- Ehsan, U.; Wintersberger, P.; Liao, Q. V.; Watkins, E. A.; Manger, C.; Daumé III, H.; Riener, A.; and Riedl, M. O. 2022. Human-Centered Explainable AI (HCXAI): Beyond Opening the Black-Box of AI. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI EA '22. New York, NY, USA: Association for Computing Machinery. ISBN 9781450391566.
- Elhage, N.; Hume, T.; Olsson, C.; Schiefer, N.; Henighan, T.; Kravec, S.; Hatfield-Dodds, Z.; Lasenby, R.; Drain, D.; Chen, C.; et al. 2022. Toy models of superposition. *arXiv preprint arXiv:2209.10652*.
- Eubanks, V. 2018. *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press.
- Ezzeddine, F. 2024. Privacy implications of explainable AI in data-driven systems. *arXiv preprint arXiv:2406.15789*.
- Farrow, R. 2023. The possibilities and limits of explicable artificial intelligence (XAI) in education: A socio-technical perspective. *Learning, Media and Technology*, 48(2): 266–279.
- Förster, M.; Broder, H. R.; Fahr, M. C.; Klier, M.; and Fink, L. 2025. Tell me more, tell me more: the impact of explanations on learning from feedback provided by Artificial Intelligence. *European Journal of Information Systems*, 1–23.
- Frank, M. P.; and George, G. 2024. Paradigm shift from AI to XAI of Society 5.0: Machine-centric to human-centric. In *XAI Based Intelligent Systems for Society 5.0*, 3–28. Elsevier.
- Fricker, M. 2017. Evolving concepts of epistemic injustice. In *The Routledge handbook of epistemic injustice*, 53–60. Routledge.
- Friedman, J. H. 2001. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, 1189–1232.
- Geiger, A.; Lu, H.; Icard, T.; and Potts, C. 2021. Causal abstractions of neural networks. *Advances in Neural Information Processing Systems*, 34: 9574–9586.
- Gerlings, J.; Jensen, M. S.; and Shollo, A. 2021. Explainable AI, but explainable to whom? An exploratory case study of xAI in healthcare. In *Handbook of Artificial Intelligence in Healthcare: Vol 2: Practicalities and Prospects*, 169–198. Springer.
- 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.
- Greenwell, B. M. 2017. pdp: An R package for constructing partial dependence plots. *The R Journal*, 9.
- Guidotti, R.; Monreale, A.; Ruggieri, S.; Turini, F.; Giannotti, F.; and Pedreschi, D. 2018. A Survey of Methods for Explaining Black Box Models. *ACM Comput. Surv.*, 51(5).
- Gummadi, A. N.; Napier, J. C.; and Abdallah, M. 2024. XAI-IoT: an explainable AI framework for enhancing anomaly detection in IoT systems. *IEEE Access*.
- Habibullah, K. M.; et al. 2024. Explainable ai: A diverse stakeholder perspective. In *2024 IEEE 32nd International Requirements Engineering Conference (RE)*, 494–495. IEEE.
- Hashmati, N.; Wörnberg, H.; Brorsson, E.; and Obaid, M. 2024. eXplainable AI Interfaces With (and for) Expert Operators: A Participatory Design Approach. In *Proceedings of the 36th Australian Conference on Human-Computer Interaction, OzCHI, Brisbane, Australia*, volume 30.
- Hoffman, R. R.; Mueller, S. T.; Klein, G.; and Litman, J. 2023. Measures for explainable AI: Explanation goodness, user satisfaction, mental models, curiosity, trust, and human-AI performance. *Frontiers in Computer Science*, 5: 1096257.
- Hohman, F.; Head, A.; Caruana, R.; DeLine, R.; and Drucker, S. M. 2019. Gamut: A design probe to understand how data scientists understand machine learning models. In *Proceedings of the 2019 CHI conference on human factors in computing systems*, 1–13.
- Horsky, J. 2025. Cognitive Support for Decisions in the Context of Clinical Workflows. In *Reengineering Clinical Workflow in the Digital and AI Era: Toward Safer and More Efficient Care*, 9–33. Springer.
- Huskonen, P.; and Korteniemi, A. 1992. Explanation based on contexts. In *Proceedings Eighth Conference on Artificial Intelligence for Applications*, 179,180,181,182,183,184,185. Los Alamitos, CA, USA: IEEE Computer Society.

- Jacovi, A.; Bastings, J.; Gehrmann, S.; Goldberg, Y.; and Filippova, K. 2023. Diagnosing ai explanation methods with folk concepts of behavior. *Journal of Artificial Intelligence Research*, 78: 459–489.
- Jain, S.; and Wallace, B. C. 2019. Attention is not Explanation. In Burstein, J.; Doran, C.; and Solorio, T., eds., *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 3543–3556. Minneapolis, Minnesota: Association for Computational Linguistics.
- Jakobsen, T. S. T.; Cabello, L.; and Søggaard, A. 2023. Being Right for Whose Right Reasons? In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1033–1054.
- Jean-Quartier, C.; Bein, K.; Hejny, L.; Hofer, E.; Holzinger, A.; and Jeanquartier, F. 2023. The Cost of Understanding—XAI Algorithms towards Sustainable ML in the View of Computational Cost. *Computation*, 11(5): 92.
- Jesus, S.; Belém, C.; Balayan, V.; Bento, J.; Saleiro, P.; Bizarro, P.; and Gama, J. 2021. How can I choose an explainer? An application-grounded evaluation of post-hoc explanations. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, 805–815.
- Kaur, H.; Conrad, M. R.; Rule, D.; Lampe, C.; and Gilbert, E. 2024. Interpretability Gone Bad: The Role of Bounded Rationality in How Practitioners Understand Machine Learning. *Proceedings of the ACM on Human-Computer Interaction*, 8(CSCW1): 1–34.
- Kaur, H.; Nori, H.; Jenkins, S.; Caruana, R.; Wallach, H.; and Wortman Vaughan, J. 2020. Interpreting interpretability: understanding data scientists’ use of interpretability tools for machine learning. In *Proceedings of the 2020 CHI conference on human factors in computing systems*, 1–14.
- Kim, B.; Park, J.; and Suh, J. 2020. Transparency and accountability in AI decision support: Explaining and visualizing convolutional neural networks for text information. *Decision Support Systems*, 134: 113302.
- Kim, S.; Choo, S.; Park, D.; Park, H.; Nam, C. S.; Jung, J.-Y.; and Lee, S. 2023. Designing an XAI interface for BCI experts: A contextual design for pragmatic explanation interface based on domain knowledge in a specific context. *International Journal of Human-Computer Studies*, 174: 103009.
- Kulesza, T.; Stumpf, S.; Burnett, M.; Yang, S.; Kwan, I.; and Wong, W.-K. 2013. Too much, too little, or just right? Ways explanations impact end users’ mental models. In *2013 IEEE Symposium on visual languages and human centric computing*, 3–10. IEEE.
- Kunz, J.; and Kuhlmann, M. 2024. Properties and Challenges of LLM-Generated Explanations. In *Proceedings of the Third Workshop on Bridging Human–Computer Interaction and Natural Language Processing*, 13–27.
- Langton, R. 2010. Miranda Fricker Epistemic Injustice: Power and the Ethics of Knowing. Oxford, Oxford University Press, 2007. *Hypatia*, 25(2): 459–464.
- Lazar, S. 2024. Legitimacy, authority, and democratic. *Oxford Studies in Political Philosophy*, 10: 28.
- Lechterman, T. M. 2022. *The concept of accountability in AI ethics and governance*. Oxford University Press Oxford, 164–182.
- Lei, T.; Barzilay, R.; and Jaakkola, T. 2016. Rationalizing Neural Predictions. In Su, J.; Duh, K.; and Carreras, X., eds., *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 107–117. Austin, Texas: Association for Computational Linguistics.
- Letrache, K.; and Ramdani, M. 2023. Explainable Artificial Intelligence: A Review and Case Study on Model-Agnostic Methods. In *2023 14th International Conference on Intelligent Systems: Theories and Applications (SITA)*, 1–8. IEEE.
- 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. New York, NY, USA: Association for Computing Machinery. ISBN 9781450367080.
- Liao, Q. V.; Zhang, Y.; Luss, R.; Doshi-Velez, F.; and Dhurandhar, A. 2022. Connecting algorithmic research and usage contexts: a perspective of contextualized evaluation for explainable AI. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 10, 147–159.
- Loyola-Gonzalez, O. 2019. Black-box vs. white-box: Understanding their advantages and weaknesses from a practical point of view. *IEEE access*, 7: 154096–154113.
- Luccioni, A. S.; Pistilli, G.; Sefala, R.; and Moorosi, N. 2025. Bridging the Gap: Integrating Ethics and Environmental Sustainability in AI Research and Practice. *arXiv preprint arXiv:2504.00797*.
- Luccioni, S.; Jernite, Y.; and Strubell, E. 2024. Power hungry processing: Watts driving the cost of ai deployment? In *Proceedings of the 2024 ACM conference on fairness, accountability, and transparency*, 85–99.
- Lundberg, S. M.; and Lee, S.-I. 2017. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30.
- Lyu, Q.; Apidianaki, M.; and Callison-Burch, C. 2024. Towards faithful model explanation in nlp: A survey. *Computational Linguistics*, 50(2): 657–723.
- Madsen, A.; Reddy, S.; and Chandar, S. 2022. Post-hoc interpretability for neural nlp: A survey. *ACM Computing Surveys*, 55(8): 1–42.
- Mansi, G.; Karusala, N.; and Riedl, M. 2025. Legally-Informed Explainable AI. *arXiv preprint arXiv:2504.10708*.
- Marasovic, A.; Beltagy, I.; Downey, D.; and Peters, M. 2022. Few-Shot Self-Rationalization with Natural Language Prompts. In Carpuat, M.; de Marneffe, M.-C.; and Meza Ruiz, I. V., eds., *Findings of the Association for Computational Linguistics: NAACL 2022*, 410–424. Seattle, United States: Association for Computational Linguistics.
- Massidda, R.; Geiger, A.; Icard, T.; and Bacciu, D. 2023. Causal abstraction with soft interventions. In *Conference on Causal Learning and Reasoning*, 68–87. PMLR.

- McDermid, J. A.; Jia, Y.; Porter, Z.; and Habli, I. 2021. Artificial intelligence explainability: the technical and ethical dimensions. *Philosophical Transactions of the Royal Society A*, 379(2207): 20200363.
- Miller, T. 2019. Explanation in artificial intelligence: Insights from the social sciences. *Artificial intelligence*, 267: 1–38.
- Mohseni, S.; Zarei, N.; and Ragan, E. D. 2021. A multidisciplinary survey and framework for design and evaluation of explainable AI systems. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 11(3-4): 1–45.
- Mökander, J.; and Floridi, L. 2022. From algorithmic accountability to digital governance. *Nature Machine Intelligence*, 4(6): 508–509.
- Mökander, J.; Juneja, P.; Watson, D. S.; and Floridi, L. 2022. The US Algorithmic Accountability Act of 2022 vs. The EU Artificial Intelligence Act: what can they learn from each other? *Minds and Machines*, 32(4): 751–758.
- Molnar, C.; König, G.; Herbringer, J.; Freiesleben, T.; Dandl, S.; Scholbeck, C. A.; Casalicchio, G.; Grosse-Wentrup, M.; and Bischl, B. 2020. General pitfalls of model-agnostic interpretation methods for machine learning models. In *International Workshop on Extending Explainable AI Beyond Deep Models and Classifiers*, 39–68. Springer.
- Muhammad, A. E.; and Yow, K.-C. 2023. Demystifying Canada’s Artificial Intelligence and Data Act (AIDA): The good, the bad and the unclear elements. In *2023 IEEE Canadian Conference on Electrical and Computer Engineering (CCECE)*, 510–515. IEEE.
- Munn, M.; and Pitman, D. 2022. *Explainable AI for practitioners*. ” O’Reilly Media, Inc.”.
- Naiseh, M.; Al-Thani, D.; Jiang, N.; and Ali, R. 2023. How the different explanation classes impact trust calibration: The case of clinical decision support systems. *International Journal of Human-Computer Studies*, 169: 102941.
- Nanda, N.; Chan, L.; Lieberum, T.; Smith, J.; and Steinhardt, J. 2023. Progress measures for grokking via mechanistic interpretability. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*.
- Nascita, A.; Montieri, A.; Aceto, G.; Ciuonzo, D.; Persico, V.; and Pescapé, A. 2021. XAI meets mobile traffic classification: Understanding and improving multimodal deep learning architectures. *IEEE Transactions on Network and Service Management*, 18(4): 4225–4246.
- Nauta, M.; Trienes, J.; Pathak, S.; Nguyen, E.; Peters, M.; Schmitt, Y.; Schlötterer, J.; Van Keulen, M.; and Seifert, C. 2023. From anecdotal evidence to quantitative evaluation methods: A systematic review on evaluating explainable ai. *ACM Computing Surveys*, 55(13s): 1–42.
- Novelli, C.; Taddeo, M.; and Floridi, L. 2024. Accountability in artificial intelligence: what it is and how it works. *AI & Society*, 39(4): 1871–1882.
- Olah, C.; Cammarata, N.; Schubert, L.; Goh, G.; Petrov, M.; and Carter, S. 2020. Zoom in: An introduction to circuits. *Distill*, 5(3): e00024–001.
- Oldenburg, N.; and Sjøgaard, A. 2025. Navigating the Informativeness-Compression Trade-Off in XAI. *AI and Ethics*.
- Rane, N.; Choudhary, S.; and Rane, J. 2023. Explainable Artificial Intelligence (XAI) approaches for transparency and accountability in financial decision-making. *Available at SSRN 4640316*.
- Retzlaff, C. O.; Angerschmid, A.; Saranti, A.; Schneeberger, D.; Roettger, R.; Mueller, H.; and Holzinger, A. 2024. Post-hoc vs ante-hoc explanations: xAI design guidelines for data scientists. *Cognitive Systems Research*, 86: 101243.
- 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.
- Roshan, K.; and Zafar, A. 2022. Using kernel shap xai method to optimize the network anomaly detection model. In *2022 9th International Conference on Computing for Sustainable Global Development (INDIACom)*, 74–80. IEEE.
- Sandu, M. G.; and Trausan-Matu, S. 2022. Comparing model-agnostic and model-specific XAI methods in Natural Language Processing. In *RoCHI*, 115–118.
- Saphra, N.; and Wiegrefe, S. 2024. Mechanistic? In Belinkov, Y.; Kim, N.; Jumelet, J.; Mohebbi, H.; Mueller, A.; and Chen, H., eds., *Proceedings of the 7th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP*, 480–498. Miami, Florida, US: Association for Computational Linguistics.
- Saraswat, D.; Bhattacharya, P.; Verma, A.; Prasad, V. K.; Tanwar, S.; Sharma, G.; Bokoro, P. N.; and Sharma, R. 2022. Explainable AI for healthcare 5.0: opportunities and challenges. *IEEE Access*, 10: 84486–84517.
- Scharowski, N.; Perrig, S. A.; Svab, M.; Opwis, K.; and Brühlmann, F. 2023. Exploring the effects of human-centered AI explanations on trust and reliance. *Frontiers in Computer Science*, 5: 1151150.
- Selbst, A. D.; Boyd, D.; Friedler, S. A.; Venkatasubramanian, S.; and Vertesi, J. 2019. Fairness and abstraction in sociotechnical systems. In *Proceedings of the conference on fairness, accountability, and transparency*, 59–68.
- Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2020. Grad-CAM: visual explanations from deep networks via gradient-based localization. *International journal of computer vision*, 128: 336–359.
- Sharkey, L.; Braun, D.; and Millidge, B. 2022. Taking features out of superposition with sparse autoencoders. In *AI Alignment Forum*, volume 8, 15–16.
- Smart, A.; and Kasirzadeh, A. 2024. Beyond model interpretability: socio-structural explanations in machine learning. *AI & SOCIETY*, 1–9.
- Smith-Renner, A.; Fan, R.; Birchfield, M.; Wu, T.; Boyd-Graber, J.; Weld, D. S.; and Findlater, L. 2020. No explainability without accountability: An empirical study of explanations and feedback in interactive ml. In *Proceedings of the 2020 chi conference on human factors in computing systems*, 1–13.

- Sokol, K.; and Vogt, J. E. 2024. What Does Evaluation of Explainable Artificial Intelligence Actually Tell Us? A Case for Compositional and Contextual Validation of XAI Building Blocks. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, CHI EA '24. New York, NY, USA: Association for Computing Machinery. ISBN 9798400703317.
- Sovrano, F.; and Vitali, F. 2021. From philosophy to interfaces: an explanatory method and a tool inspired by achinstein's theory of explanation. In *Proceedings of the 26th International Conference on Intelligent User Interfaces*, 81–91.
- Spartalis, C. N.; Semertzidis, T.; and Daras, P. 2023. Balancing xai with privacy and security considerations. In *European Symposium on Research in Computer Security*, 111–124. Springer.
- Stojkovic, J.; Choukse, E.; Zhang, C.; Goiri, I.; and Torrellas, J. 2024. Towards greener llms: Bringing energy-efficiency to the forefront of llm inference. *arXiv preprint arXiv:2403.20306*.
- Stone, J. I.; and Mittelstadt, B. 2025. Legitimate power, illegitimate automation: The problem of ignoring legitimacy in automated decision systems. *ACM Journal on Responsible Computing*.
- Sundararajan, M.; Taly, A.; and Yan, Q. 2017. Axiomatic attribution for deep networks. In *International conference on machine learning*, 3319–3328. PMLR.
- Vale, D.; El-Sharif, A.; and Ali, M. 2022. Explainable artificial intelligence (XAI) post-hoc explainability methods: Risks and limitations in non-discrimination law. *AI and Ethics*, 2(4): 815–826.
- Veale, M.; and Binns, R. 2017. Fairer machine learning in the real world: Mitigating discrimination without collecting sensitive data. *Big Data & Society*, 4(2): 2053951717743530.
- Veale, M.; and Brass, I. 2019. Administration by algorithm? Public management meets public sector machine learning. Oxford University Press.
- Veale, M.; and Zuiderveen Borgesius, F. 2021. Demystifying the Draft EU Artificial Intelligence Act—Analysing the good, the bad, and the unclear elements of the proposed approach. *Computer Law Review International*, 22(4): 97–112.
- Vilone, G.; and Longo, L. 2021. Notions of explainability and evaluation approaches for explainable artificial intelligence. *Information Fusion*, 76: 89–106.
- Wang, D.; Yang, Q.; Abdul, A.; and Lim, B. Y. 2019. Designing theory-driven user-centric explainable AI. In *Proceedings of the 2019 CHI conference on human factors in computing systems*, 1–15.
- Wang, Z.; Chu, Z.; Doan, T. V.; Ni, S.; Yang, M.; and Zhang, W. 2024. History, development, and principles of large language models: an introductory survey. *AI and Ethics*, 1–17.
- Wexler, J.; Pushkarna, M.; Bolukbasi, T.; Wattenberg, M.; Viégas, F.; and Wilson, J. 2020. The What-If Tool: Interactive Probing of Machine Learning Models. *IEEE Transactions on Visualization and Computer Graphics*, 26(1): 56–65.
- Wiegrefe, S.; Hessel, J.; Swayamdipta, S.; Riedl, M.; and Choi, Y. 2022. Reframing Human-AI Collaboration for Generating Free-Text Explanations. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 632–658.
- Wolf, C. T. 2019. Explainability scenarios: towards scenario-based XAI design. In *Proceedings of the 24th International Conference on Intelligent User Interfaces*, 252–257.
- Wolf, C. T.; and Ringland, K. E. 2020. Designing accessible, explainable AI (XAI) experiences. *ACM SIGACCESS Accessibility and Computing*, (125): 1–1.
- Zhang, C. A.; Cho, S.; and Vasarhelyi, M. 2022. Explainable artificial intelligence (XAI) in auditing. *International Journal of Accounting Information Systems*, 46: 100572.
- Zhu, X.; Wang, D.; Pedrycz, W.; and Li, Z. 2022. Fuzzy rule-based local surrogate models for black-box model explanation. *IEEE Transactions on Fuzzy Systems*, 31(6): 2056–2064.