

Emergent, not Immanent: A Baradian Reading of Explainable AI

Fabio Morreale
Sony AI
Barcelona, Spain
fabio.morreale@sony.com

Joan Serrà
Sony AI
Barcelona, Spain
joan.serra@sony.com

Yuki Mitsufuji
Sony AI
New York, USA
yuki.mitsufuji@sony.com

Abstract

Explainable AI (XAI) is frequently positioned as a technical problem of revealing the inner workings of an AI model. This position is affected by unexamined onto-epistemological assumptions: meaning is treated as immanent to the model, the explainer is positioned outside the system, and a causal structure is presumed recoverable through computational techniques. In this paper, we draw on Barad’s agential realism to develop an alternative onto-epistemology of XAI. We propose that interpretations are material-discursive performances that emerge from situated entanglements of the AI model with humans, context, and the interpretative apparatus. To develop this position, we read a comprehensive set of XAI methods through agential realism and reveal the assumptions and limitations that underpin several of these methods. We then articulate the framework’s ethical dimension and propose design directions for XAI interfaces that support emergent interpretation, using a speculative text-to-music interface as a case study.

CCS Concepts

• **Human-centered computing** → **HCI theory, concepts and models**; • **Computing methodologies** → **Philosophical/theoretical foundations of artificial intelligence**.

Keywords

Interpretability, explainability, diffraction, incommensurability, transparency, agential realism

ACM Reference Format:

Fabio Morreale, Joan Serrà, and Yuki Mitsufuji. 2026. Emergent, not Immanent: A Baradian Reading of Explainable AI. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3772318.3790725>

1 Introduction

Explainable¹ AI (XAI) has been the subject of a large amount of multidisciplinary research, including HCI, for decades. Despite this work, the capacity to interpret AI systems continues to be identified

¹Related terms—in particular, *interpretability* and *transparency*—have been proposed, and different stakeholders attribute them different meanings [34]. Here, we adopt them as functionally equivalent [75], as their finer differentiation lies beyond the scope of this work.



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

CHI '26, Barcelona, Spain

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2278-3/26/04

<https://doi.org/10.1145/3772318.3790725>

as one of the most pressing and unresolved objectives by scholars [32, 108] and governments alike [58]. While work focusing on identifying its sociotechnical aspects is growing [32, 35, 89, 109], prevailing discourses tend to frame interpretability as a technical problem of revealing a causal structure [110], which can only be solved through better computational tools or mathematical functions. These discourses align with representationalist, mechanistic, and positivist conceptions of explanation, which are based on three main assumptions: i) the *explanandum* (what needs explaining) exists within the model; ii) (some) *explananda* are discoverable and accessible; and iii) the explanation tools and the human explainer have negligible epistemic effect on the *explanandum*.

While these assumptions are rarely made explicit, they underpin much of mainstream XAI and have sedimented into a vernacular image of AI models as “black boxes”. This metaphor, as Langdon Winner warned, naturalises obscurity while foreclosing inquiry into the socio-technical constitution of the system [117]. Bruno Latour similarly argued that technologies become black-boxed through deliberate acts of closure, thus opacity is not a neutral property of systems but the effect of sociotechnical work [71]. Ananny and Crawford further criticised this metaphor for sidelining the material and ideological complexities of *seeing* and for assuming that accountability can be simply achieved by exposing internal workings [5]. Clinging to this metaphor reaffirms an *ontology* in which meaning resides inside the model, prior to and independent of any constitutive practices; an *epistemology* premised on uncovering such hidden meaning; and an *ethics* that equates accountability with transparency. What is needed is an alternative ethico-onto-epistemology that rethinks from the ground-up what AI interpretations are (ontology), how they are generated and understood (epistemology), and they relate to accountability and responsibility (ethics). Several scholars have already raised critical concerns about dominant assumptions in mainstream XAI [1, 22, 34, 68, 90]. However, these contributions are not intended as onto-epistemological readings of the commitments that underpin prevailing notions of interpretation in XAI.

In this paper, we develop such an onto-epistemological reading by drawing on Karen Barad’s agential realism framework and the associated epistemic optic of *diffraction* [6], itself inspired by Donna Haraway’s work [52]. Barad first developed diffraction as a method to diagnose the problematic onto-epistemological assumptions of scientific practice and eventually developed it into a complete alternative metaphysics: agential realism [6]. Within agential realism, knowledge is not produced by a detached subject observing an independent object, but emerges through situated *intra-actions* in which the subject and the object are mutually constituted. In this view, knowing is a practice of world-making, thus epistemology and ontology are inseparable.

Agential realism and diffraction have been applied to multidisciplinary research and, increasingly, in HCI [13, 16, 47, 78, 86, 87, 92, 96, 98, 99, 103, 118] but, to the best of our knowledge, not directly to XAI. An agential-realist reading of XAI challenges the assumptions i–iii listed above, showing that explanations might, after all, not reside within the model, waiting to be discovered and extracted, but rather emerge from situated intra-actions between the AI model, human interpreters, datasets, XAI tools, and the socio-technical context through which they are performed. This reading shifts interpretability from an exercise of uncovering pre-existing explanations to a material-discursive practice in which the entanglements of more-than-human actors participate in the co-constitution of the explanation. To develop this position, we examine a set of representative XAI methods through Baradian optics. Our analysis reveals conceptual shortcomings inherent in the XAI methods that treat explanations as pre-existing entities awaiting discovery by a human observer, while also showing that some of these methods in fact operate in ways that are better described as diffractive.

The main contribution of this paper is a novel onto-epistemological reading of XAI in which interpretability is reframed as a material-discursive performance. This contribution intervenes in debates on critical XAI [1, 21, 22, 40, 48, 68, 90] by understanding interpretability through entanglement-based [41] and more-than-human [39, 42] perspectives common in HCI, which acknowledge that human and non-human actors participate in generating of knowledge and shaping reality, and recognise technologies as co-constitutive rather than neutral. While HCI has long engaged with XAI matters through a heterogeneous body of work [20, 63, 95, 115], particularly in relation to interactive and human-in-the-loop approaches to explanation [10, 11] and to responsible and socially situated perspectives on XAI [27, 33, 36, 73, 74], entanglement theories have not yet been directly brought to bear on interpretability. We made this connection explicit, thereby both advancing XAI and assessing the reach of HCI entanglement frameworks in a new setting. Also, the ethical implication of our proposed framework joins HCI discourse on ethical AI and social justice, showing how explanation practices are entangled with institutional and economic configurations, which diffractive approaches can help surface. In addition to this main contribution, this paper also offers possible design directions for XAI interfaces centred on emergent interpretations, proposing that multiple and even contrasting explanations should be staged and negotiated through situated, ambiguous interfaces.

This paper is structured as follows. Section 2 introduces the philosophical theories and concepts that ground our work. Section 3 overviews a large number of XAI methods, which are analysed using Baradian optics as epistemic tools in Section 4. The discussions in Section 5 introduce the ethical consequences of our revised understanding of XAI, and Section 6 develops design directions for XAI interfaces, illustrated through a speculative text-to-music case study. Finally, Section 7 draws conclusions, outlines study limitations and discusses future work.

2 Background

In this section, we introduce philosophical concepts and scholarly work that ground our work.

2.1 Positivism and Representationalism

Science and engineering research, including a substantial portion of current HCI work, inherits from a philosophical tradition of *positivism*, which proposes that an objective reality exists and scientific tools can be used to observe and make sense of it in a value-free manner [107]. Relatedly, *representationalism* is a description used by authors opposing the idea that reality is always already mediated through epistemic processes like signs and language, which have “the power to mirror preexisting phenomena” [6]. This formulation is grounded in the metaphysical separation of the subject and the object [23] and presupposes that there exists a “distinction between representations and that which they purport to represent”, and that individuals exist “awaiting/inviting representation” [6]. Thus, *beings* exist anterior to, and independent of, their representations, and instrumentation is separate from what they measure [57]. Beginning with second-wave feminism, and carried forward by contemporary feminist scholarship, radical critiques of Western philosophy have challenged its tenets [107], including representationalism and positivism, especially in the work of authors like Michael Foucault, Judith Butler, and, in particular, Joseph Rouse and Karen Barad.

2.2 Optics of Knowledge Production

Østerlund and colleagues [119] systematised Barad’s use of optical metaphors [6] into three distinct optics—reflective, refractive, and diffractive—offering a framework for analysing different epistemological stances in knowledge production.

Refraction was only indirectly mentioned by Barad, who grouped it with reflection [119]. However, Østerlund and colleagues suggest that the refractive optic warrants separate treatment. In this optic, knowledge is produced by revealing, through precise observations, pre-existing structures that “exist out there”. The metaphor at play is that of a *lens*, which offers direct access to stable and observable objects. A refractive approach thus treats the measurement apparatus as neutral. This stance underpins many engineering and computational traditions, and continues to dominate large portions of scientific practice.

Reflection repositions knowledge production as an interpretative act of a reality reflected by a *mirror*. This optics maintains that a reality with an intrinsic structure and pre-given objects exists and is available to an observer, but admits the impossibility of gaining direct access to reality. The observer can only capture distorted or incomplete reflections, thus the produced knowledge “is not a perfect replica but a partial reconstruction of reality” [119]. In this optic, the apparatus is no longer entirely neutral, but is still separate from the object and the observer. The world remains pre-structured; the observer is now just more aware of their limitations in representing it. The aim is thus still representational, but this optic acknowledges the interpretative work involved in making sense of the observed object. Barad criticises this framework for still *holding the world at a distance* and for not solving the epistemological gap between subject and object, while acknowledging the attempt to “put the investigative subject back in the picture” [7].

Diffraction, in physics, refers to the bending and spreading of waves when they encounter an obstacle or pass through a narrow opening. A classic demonstration of diffraction is Thomas Young’s

double-slit interference experiment from 1803. In this experiment, light passing through two closely spaced slits produces, on a screen, an interference pattern² of alternating bright and dark bands. This phenomenon is also perceivable when a razor blade is illuminated by monochromatic light: its shadow is not sharply delineated, but a pattern of alternating light and dark lines is exhibited [7]. In the twentieth century, similar experiments were repeated with individual photons and other quantum particles. In these versions, the appearance or disappearance of the interference pattern depends on whether the experimenter attempts to measure which slit the particle passes through. It was in this context that physicist Niels Bohr introduced his principle of complementarity [15], arguing that quantum entities cannot be fully captured by either a purely particle-like or a purely wave-like description: what is observed depends on the measuring apparatus. Bohr's work on complementarity provided the scientific grounds for challenging representationalism. Barad extended this insight beyond physics, mobilising it into metaphysics to articulate the concept of diffraction, which she later developed into agential realism.

2.3 Agential Realism

Agential realism is Barad's broader onto-epistemological framework, of which diffraction is one methodological optic. Agential realism posits that i) there is no absolute separation between observer and observed; ii) measurements do not *represent* "measurement-independent states of being" [6]; and iii) observations are not transparent windows onto reality. Rather, the observer and the observation/measurement apparatus are part of an entanglement that co-determines, and thus cannot be separated from, what is observed. These entities thus do not exist prior to their relations, but are configured by and materialise through their entanglement; they *intra-act* [72].

Rather than assuming interactions between already constituted entities, intra-action emphasises how phenomena take form through relational processes [54]. In this view, what classical epistemology treats as an ontological separation between an observing subject and an observed object is understood as *phenomena*, which are ontological primitive relations that do not presuppose pre-existing *relata* [7]. Intra-actions address the Cartesian subject-object distinction by producing an *agential cut*, which temporarily resolves indeterminacy and marks provisional boundaries between subject and object. Outside any specific intra-action, subject and object remain indeterminate [7]. Such cuts are enacted through *apparatuses*, which are not static instruments that measure pre-existing and measurement-independent entities but "specific material practices through which local semantic and ontological determinacy are intra-actively enacted" [7]. Apparatuses co-produce phenomena by enacting these cuts, and thus drawing boundaries in sociomaterial reality [94, 119], and are exclusionary as they make some aspects visible while rendering others invisible.

Under agential realism, knowledge is *performative*: practices of knowing participate in the creation of reality [22] and thus nothing exists before it is produced as a *discursive phenomenon*. Ontology

²The pattern arises from the superposition of the diffracted waves emerging from the slits: where crests meet, the waves reinforce each other (bright), and where a crest meets a trough, they cancel (dark).

and epistemology thus implicate each other: "there are no entities and no knowable characteristics outside of practices of knowing them" [22]. Entities like the observer, the apparatus, and the observed are not given in advance but emerge via their entanglement through material-discursive practices. The diffraction pattern offers a metaphor for this process: as in the case of a razor blade illuminated by monochromatic light, sharp boundaries dissolve into interference fringes. Boundaries are thus not fixed in advance but emerge as effects of relational configurations. As Barad writes, "a diffraction pattern does not map where differences appear, but rather maps where the effects of differences appear." The diffraction pattern is thus not a representation of what is, but a trace of what has happened and the configuration that brought the phenomenon into being [7].

2.4 Feminist Epistemology and XAI

Donna Haraway, a foundational figure in feminist science and technology studies, criticised the universality and alleged objectivity of mainstream scientific knowledge as incapable of seeing its own biases [53]. Haraway's proposed alternative is a *situated* knowledge, which is the product of specific socially-situated (e.g., historical, geographical, and cultural) circumstances. Situatedness, thus, turns down purportedly objective "views from nowhere" in order to focus on "acquiring knowledge in particular contexts" [51] by means of a dialogue between subjects in a hierarchical relationship of observer and observed, as proposed by black feminist epistemology [55].

Feminist epistemology has been previously employed to assess the assumptions embedded in AI research [19, 26, 30, 112], but relatively few works have applied feminist lenses to XAI. Hancox-Li and Kumar suggested that XAI research should allow for multiple interpretations [51] and should refocus on subjugated and marginalised populations, allowing them to participate in AI explanations, when relevant. Similarly, Nicenboim and colleagues [90] drew upon feminist and posthuman theories to criticise the very idea of explainability. They suggest that a more apt concept is that of *shared understanding* between humans and non-humans involved in explanations. They also reject the idea of a single, correct explanation and invite to explore multiple co-produced understandings.

A recent work on XAI for bias reduction borrows concepts from feminist epistemology in critiquing what the authors term *technical XAI* [59]. They suggest that XAI models can only detect a limited amount of bias because they are only simplified representations of real AI models, and thus "can detect only those biases that their simplifying assumptions are attuned to". The authors also argued against the objectivity of XAI models, as they are affected by undiscussed assumptions and unquestioned social values. Finally, Klumbyte and colleagues [64, 65] argued that feminist principles can enrich XAI by challenging universal explanations, foregrounding power and accountability, and centring marginalised perspectives through participatory and interactive approaches. Despite these contributions, diffraction and agential realism have not yet been systematically brought to bear on XAI. Conversely, research in HCI has increasingly drawn on these theories.

2.5 Agential Realism and Diffraction in HCI

Agential realism and diffraction have primarily informed broader strands of theoretical HCI and, in particular, entanglement HCI [41], posthuman and more-than-human perspectives [39], and sociomaterial approaches [94]. Barad's work has also been applied in practical ways to a variety of domains, including design research [13], choreography [16], lived data [99], environmental data [47], wellbeing technologies [118] and, particularly, interactions with digital musical interfaces [78, 86, 87, 96, 98, 103]. Across these works, agential realism and diffraction have been used to better understand aspects like authorship [16], ambiguity [96], unpredictability [87], creativity [13], and agency [86] in more-than-human interactions, as well as to surface assumptions about what sort of "humans" do specific technologies embed [118].

Furthermore, diffraction has informed design and research methods. Giaccardi and colleagues proposed *diffractive interfaces* in the context of forest simulation data [47]. They showed how agent-based models enact agential cuts, privileging particular attributes and measurements, and thereby reducing the more-than-human complexity of forests. Their proposed diffractive interfaces make these cuts explicit and enable stakeholders to move between different cuts and measurement practices, exploring how different ways of measuring and representing forest data bring different realities into being. Also inspired by Barad is Morrison and McPherson's method of *diffractive dialogue* [86], in which each author responded to a series of discussion prompts by drawing on their own disciplinary background.

Despite this growing engagement with Baradian work, to the best of our knowledge, our work is the first to develop a systematic Baradian onto-epistemology of more-than-human configurations in XAI.

3 XAI Methods

In this section, we review common XAI methods, adapting and complementing the taxonomy proposed by [102] with additional techniques relevant to our analysis. Our aim is not to provide a systematic survey, but to offer an overview of representative approaches.

3.1 Feature Attribution

Feature attribution methods typically aim to explain a *single* prediction generated by a neural network by identifying which (input or intermediate) features were most influential for an output. The assumption is that the model's decisions are determined by certain features, the influence of which is inherent in the model, and that an observer should be able to discover (thereby revealing), at least in part, the model's internal decision logic.

Saliency maps aim to discover the pixels (or regions) that strongly affect an image classification task and, thus, show where the model "looked" in the input. This technique belongs to visualisation-based explainers, which are considered "objective" explanation tools since they rely on the model's own signals (e.g., gradient values), rather than human input [75]. A typical example in computer vision is Grad-CAM [104], which aims to highlight the regions of an image that are most important for a specific prediction.

Perturbation-based approaches systematically modify the input (e.g., masking it, adding noise, shifting values) and observe the effect of these modifications on the output. **Ablation-based** approaches are closely related, but intervene instead within the model itself, removing specific components (e.g., neurons, layers) in order to assess their relative contribution to the model's predictions.

Surrogate models are simplified models specifically designed for XAI, which approximate the behaviour of more complex models. LIME [97], a classical example of a surrogate model, constructs local, interpretable models by approximating the original model's behaviour using a small linear model. It assumes that complex models are locally linear, allowing the resulting weights to explain the factors influencing the output. Another example is SHAP.³ SHAP values are reportedly the *unique* solution that satisfies a set of desirable axioms for feature importance and are thus believed to yield a single objectively "correct" attribution of credit to features, implying that a model's prediction can be decomposed exactly into feature contributions.

Decomposition-based methods break down a model's reasoning process and attribute its predictions to specific input features or intermediate computations. Layer-wise relevance propagation (LRP) is one such method [81], designed to explain which input features support a model's outputs by breaking it down into relevance scores and propagating them backwards through the network. For example, when classifying an image of a dog, LRP decomposes the model's "dog" prediction into pixel-level relevance scores, revealing that high-relevance regions (e.g., the dog's ears and nose) most strongly supported the output. The output is a relevance heatmap, which provides pixel/neuron-level scores but no human-readable concepts.

3.2 Sample-based

Sample-based methods aim to explain model behaviour by analysing input-output relationships across *multiple* examples. Rather than attributing importance to individual features of a single instance, these methods characterise predictions in relation to several samples.

Adversarial example [50] methods perform minimal input perturbations that flip outputs while appearing unchanged to humans, thereby possibly revealing what features are critical for decision-making. The central idea is that if a minimal change can flip a prediction, the modified parts are likely important for that prediction.

Counterfactual explanations identify minimal and semantically meaningful changes to the input that alter the model's output. Rather than unveiling internal model structure, counterfactuals focus on how the outcome would change under small changes to the input [113]. The objective is not to open the model's black box, but to show how different inputs would lead to different outcomes. The perturbation must be noticeable to a human: for example, the goal might be to modify an input so that an 'entry denied' outcome becomes 'entry approved'.

³The taxonomy we follow [102] includes SHAP [77] among the surrogate models. SHAP, however, is rather an explanation model for a model's behaviour (or an approximation of it), but it does not replace the original model with a surrogate.

Contrastive explanations address the question of why a model produced a specific output instead of a plausible alternative [79]. Rather than justifying the prediction in isolation, this method explains the decision by identifying the critical, discriminative features that differentiate the input from cases that would have resulted in the alternative.

Prototype-based methods such as ProtoPNet [25] explain predictions by attributing them to explicitly learned *prototypes* (e.g., prototypical image patches) and highlighting the corresponding regions in the input. Explanations are generated by visualising the input patches that activate the most similar prototypes (e.g., “the network classifies this as a clay-colored sparrow because the wing pattern resembles that of this prototype sparrow”). By making its reasoning process dependent on this direct comparison to prototypical parts, ProtoPNet aims to align its decision process with human conceptualisation.

Influence functions techniques are used to estimate which training examples had the greatest impact on a prediction [66]. Instead of re-running the original learning algorithm, the method approximates how the model’s behaviour would change if the weight of a training point was emphasised.

3.3 Probing-based

While methods defined in 3.1 and 3.2 are applicable to inputs and outputs, probing-based methods aim to inspect the internal model representations to understand whether they encode specific knowledge. To do so, auxiliary simple classifiers (probes) are trained on those representations to predict interpretable, pre-determined features.

Knowledge probing investigates the linguistic and factual knowledge implicitly encoded within a model’s internal representations. Probes are trained on the model’s embeddings to predict specific properties, such as syntactic or semantic labels. For example, BERT [29] has been probed to see if information like part-of-speech tags could be recovered directly from its hidden states [111]. A probe’s high accuracy on a task is interpreted as evidence that the corresponding knowledge is stored at that particular layer of the model.

Concept-based probing starts from human-defined concepts (e.g., colour, gender, sentiment) and tries to map them to *directions* in the embedding space. A key method is TCAV (Testing with Concept Activation Vectors) [61], which quantifies the influence of a concept on a model’s predictions. In TCAV, a user defines a concept (e.g., stripes) by providing a set of exemplary inputs that contain it (e.g., images of various striped patterns). TCAV learns a direction in the activation space that represents that concept and measures the model’s sensitivity to changes along it. For instance, it can quantify how much the prediction *zebra* depends on the concept *stripes*.

Neuron activation-based probing operates on the premise that the activations of individual (or groups of) neurons correspond to interpretable, often human-defined concepts. A prominent example [12] employs a larger model (GPT-4) to automatically generate interpretations of the patterns that activate specific neurons in a model. By analysing the responses of neurons across diverse inputs, this approach aims to map neuronal function to discrete elements of human understanding.

3.4 Mechanistic Interpretability

Mechanistic interpretability methods aim to “understand neural networks’ decision-making processes” [106] and are based on two main assumptions: i) that complex AI systems have similar ontological status to biological organisms, and thus require similar epistemic tools to be investigated [70]; ii) that complex network behaviours can be explained by identifying a simpler algorithm and a sparse subset of the network that can execute it [88]. Two main approaches in mechanistic interpretability can be identified: reverse-engineering and circuits [106].

Reverse engineering approaches decompose the network into components and then attempt to identify their functions, hoping that structures in the hidden activations correspond to those of the neural computation [106]. The most intuitive candidates for decomposition, individual neurons and attention heads, turned out not to be suitable, as they are not a “natural unit for human understanding” [18] and are *polysemantic*; i.e., they encode more than one meaning. Proponents of this method suggested that polysemanticity might be caused by *superposition*, a hypothesised phenomenon where a neural network represents more independent features than it has neurons or activations.

Researchers sought to address superposition by finding *sparse* decompositions [3, 18]; i.e. by decomposing activations into sparsely active features that might correspond to human-interpretable concepts [3]. The subsequent stage involves formulating and testing hypotheses about the functional roles of these features and their interactions [106]. To map these interactions, researchers use tools like *attribution graphs* to visualise computational pathways and causal relations to highlight patterns thought to be most relevant for explaining network behaviour. To manage complexity, these graphs are typically simplified through pruning.

The **circuits** method focuses on identifying interpretable sub-graphs within a network that appear to implement specific algorithms, behaviours, or capabilities [93]. To address the problem of polysemantic neurons, circuits construct a local replacement model in which neurons are replaced with interpretable building blocks (*features*), which are intended to be monosemantic, producing a surrogate subnetwork that approximates the original model [3]. These replacement models are then studied to trace how outputs are computed. This approach presumes a correspondence between internal structures (which are not metaphorical but are intended as faithful implementations of cognitive subroutines) and meaningful behaviour, treating neural networks as systems whose logic can be rendered in human-legible terms.

3.5 Interactive XAI

Interactive XAI does not refer to a specific method, but to a way of engaging users interpreting XAI models. Rather than receiving static outputs, humans actively shape the explanatory process. These approaches treat interpretability as a dialogue, enabling users to interrogate models, test hypotheses, and refine explanations through iterative interaction and direct manipulation.

Interactive explanations allow users to probe, steer, and co-construct the criteria for a satisfactory explanation. Research in HCI underscored that interaction significantly enhances model

legibility. This is exemplified by systems like Gamut, a visual analytics tool designed to investigate how interactive interfaces can enhance model interpretation [56]. A complementary paradigm is **explanatory debugging**, which is based on a collaborative loop: the system details the reasoning behind its predictions, and the user, in turn, provides corrective feedback, making explanation a continuous dialogue rather than a fixed artefact [69]. Knowledge about how the model works materialises gradually, through the user and the system in tandem, refining it.

Interactive counterfactual explanations allow users to modify an input instance (e.g., an image region or text segment), and observe how the model’s prediction and accompanying explanation change in response. An example is the What-If Tool [116], which provides an interactive interface for probing models by altering input features and immediately observing the impact on prediction.

Interactive feature exploration tools provide direct control over explanatory content through a GUI, as in Prospector [67] or ModelTracker [4]. Users can filter or highlight features based on importance, reducing cognitive load. Similarly, **interactive detail-on-demand overlays** can reveal additional context (e.g., labelling an image region) only upon user interaction (e.g., hovering or clicking), avoiding upfront information overload.

Finally, **interactive debugging tools** like AGDebugger [37] allow developers to reset an AI agent to a prior state, effectively forking its execution path to test alternative plans. This supports direct manipulation of agent behaviour, enabling developers to interactively formulate and test hypotheses about agent reasoning under varying instructions.

4 A Baradian Analysis of XAI methods

In this section, we deploy the Baradian optics of refraction, reflection, and diffraction (with the broader agential realist framework) to interrogate the onto-epistemological assumptions embedded in XAI methods. These optics are research probes rather than classificatory tools, as our aim is not to rank methods, but to expose distinct facets of their philosophical commitments. To support our reasoning, we also represented each optic in mathematical formulation and with pictorials, which we clustered in Table 1, to provide complementary registers to engage with our propositions.

4.1 Refractive and Reflective Reading of XAI

In most XAI literature, interpretability is simply treated as an intuitive *desideratum*, but left undefined. In contrast, Erasmus and Brunet characterise interpretation as “something that one does to an explanation to make it more understandable” and built a theory of interpretability on top of this definition [38]. Here, “one” is usually a human expert and the “doing” presumes that an explanation already exists, awaiting clarification.

4.1.1 Searching Immanent Explananda. Many XAI methods adopt this stance, treating AI models as if they contain a knowable (discoverable and accessible) essence that they can reveal. This approach evokes an archaeological work, where the task is to uncover structures that are presumed to lie dormant within the model. *Grad-CAM*, for instance, considers that relevance pre-exists in internal feature maps, which are presented as evidence of where the concept (already) resides in the network’s computations. *Circuits’* proponents

also aim to recover pre-existing structures and internal mechanisms, and relate to features and connections as real components that “can be rigorously studied and understood” [93]. Mechanistic interpretability researchers indeed often frame their work as akin to scientists uncovering pre-existing “true statements” about reality, with frequent analogies to biology and neuroscience [3, 76]. *Knowledge probing* tools aim at uncovering some latent linguistic knowledge and thus assume that neurons “have meanings” latent within them. As another example, *ProtoPNet’s* architecture is designed to use human-understandable reasoning (prototypes) in its logic, thus assuming that meaning is entirely communicable and interpretable via a set of analogies that humans can relate to.

These methods exemplify refractive and reflective onto-epistemologies as they assume the existence, in AI models, of knowable concepts that exist prior to and outside of the act of interpretation, waiting to be uncovered. We call this an *immanent*⁴ *ontology of XAI*: human-legible *explananda* are already present in the model, and explanation is a technical matter of directly accessing (reflective) or approximating (refractive) them. While they both see the *explananda* as pre-existing interpretations, they differ in how such concepts reach the human observer.

The **refractive** reading treats XAI tools as lenses that project sharp images of actual content (Table 1a). *Saliency maps*, for instance, work on the assumption that they can surface “the most meaningful parts” of the image [62], thus meaning was in the model all along, and the method merely isolates and displays it. *LRP* also sees explanation as an *exact decomposition* of a model’s output onto input features. The assumption is that the model’s reasoning for a given image exists as a concrete distribution of relevance over the pixels, which LRP neatly uncovers, with no distortion (Eq. 2). Similarly, *influence functions* assume that predictions are decomposable into quantifiable shares of responsibility assigned to individual training samples. Yet, the predicted explanation is not a causal truth, but a narrative produced through mathematical approximations of model behaviour treated as an objective trace of the model’s computation that the human simply reads.

The **reflective** reading recognises that XAI methods only approximate pre-existing *explananda* (Table 1b). Thus, in its initial stage, the reflective apparatus operates like a mirror that renders partial or distorted reflections, or approximations, of the AI model (Eqs. 3 and 4). For instance, *SHAP* values provide a simplified account of feature contributions that flattens the actual model’s reasoning. Similarly, *perturbation-based* methods approximate feature relevance by selectively masking or altering inputs that supposedly leave the model unchanged. The explanations they produce are *post-hoc* estimates of feature importance, and, as such, they are a map of potential decision pathways, not a transparent window into the model’s fixed internal computations.

4.1.2 Approximating Approximations. Model approximation is especially pronounced in *ante-hoc* methods as they dismember, simplify, and then recombine the model with human-palatable approximations. Here, the (still pre-existing) object is actively reshaped by


⁴We prefer the term *immanent* over *inherent*, which is more commonly used in XAI literature, since it captures the assumption that explanations are thought to reside within the model itself, as an inner essence waiting to be disclosed, rather than simply being a technical property built into its design.

(a) Refraction

$$M' = \phi_A(M), \quad \phi_A \approx I \quad (1)$$

$$E = \psi_H(M') \quad (2)$$

M' is the representation of the model M provided by the apparatus A (Eq. 1), which is assumed transparent ($\phi_A \approx I$). E is the explanation produced by the human H interpreting M' (Eq. 2).




(b) Reflection

$$\tilde{M} = \phi_A(M, A), \quad \phi_A \neq I \quad (3)$$

$$\tilde{E} = \psi_H(\tilde{M}) \quad (4)$$

\tilde{M} is a reflected/distorted version of M produced by the apparatus A (Eq. 3). \tilde{E} differs from an ideal explanation E as the apparatus is not transparent ($\phi_A \neq I$) and the interpretation ψ is affected by polarised human optics (Eq. 4).



(c) Diffraction / Agential realism

$$\delta_i = \mathcal{D}(M, A_i, H_i, C_i) \quad (5)$$

$$\Delta = \sum_{i \in \chi} \delta_i \quad (6)$$

$$E = \Psi_H(\Delta) \quad (7)$$

$$H' = \rho(\Delta, H), \quad H' \neq H \quad (8)$$

\mathcal{D} is the entangling operation generating configurations δ_i from the intra-action of M , A_i , H_i , and different contexts C_i (Eq. 5). The interference pattern Δ arises as the structured composition of different δ_i under the agential cut χ (Eq. 6). The human H interprets Δ to produce the explanation E (Eq. 7) and dynamically evolves or adapts (Eq. 8).




Table 1: Comparison of the three Baradian optics applied to XAI.

the human interpreter prior to its discovery. The interpreter thus not merely observes, but directly *intervenes* in the object of study. Examples include *LIME*, which approximates the *explanandum* with a local surrogate model, and concept erasure models like *LEACE* [9], which remove information correlated with certain concepts. In *circuits*, the artificial neurons are replaced with the local replacement model, which is composed of (supposedly interpretable) approximations. These approximations are “imperfect labels” [3] and the

model is known to differ from the original model [3]. Similarly, *attribution graphs* are pruned to only retain nodes and edges deemed to significantly contribute to the output [76].

At the successive stage, an interpretation is required to give meaning to these approximations. The reflected models are here subjected to further distortion when researchers, in the act of interpreting them, impose onto them subjective human optics (Eq. 4). *Circuits* researchers, for example, group related features into *super-nodes* [3] by identifying human-detectable patterns. These patterns, however, are approximations of approximations and, as such,

Feature 3 top activations
 al stunt rider." "* MUSIC:" " Indiana Jones Theme" "I'm not really a 9 to 5 off:
 are you supposed to be?" "I'm Indiana Jones..." "Professor of archeology, expert
 ." "My mom said I looked like Indiana Jones" "Your mom loved you very much." "T
 ie now." "Bye-bye." "This is Indiana Jones not Star Trek." "There should be no
 . Ito." "Fine!" "I'm going as Indiana Jones ." "Then give me backt Deidre's wig,
 as thinking of dressing up as Indiana Jones's mocha-skinned love child." " India

Feature 7 top activations
 d the red and black jacket in Michael Jackson's "Thriller" video which I've nev
 the "Billie Jean" video," "Michael Jackson wanted something cinematic with --
 vies." "I mean you dance like Michael Jackson and have adopted Indian superstar
 when I say, "bad," "I mean, Michael Jackson Bad." "[upbeat music]" " Thank you,
 ." "Even more aggressive than Michael Jackson's treatment of his own groin on sta
 ee." "The one who dances like Michael Jackson" ""Introduce me to him, sir." "I'

Feature 9 top activations
 -andrewl= Wall timers for my air conditioners. I have window units in my bedroom
 und can handle the loa of an air conditioner. Now the AC turns on in the middle
 that I'd have to leave both air conditioners on all day just so the place was l
 r ever since I finally got an air conditioner last week. ToyKeeper> ... but cool
 look out." "Hey, wanna buy an air conditioner?" "No, thanks." "Sit still, please.
 esent invention relates to an air conditioner for a vehicle that is used to condu

Figure 1: A few of the features grouped into the "Michael Jordan" supernode in [3].

are prone to errors. For example, one such supernode from [3] comprises features that researchers believe activate for "Michael Jordan"; yet, some of these features show stronger activation for Indiana Jones, Michael Jackson, and air conditioners than for Michael Jordan (see Fig. 1).

The XAI models that operate within these two optics are constructed to identify human-readable and observer-independent properties (e.g., concepts, mechanisms) in a model, even when it has none. Mechanistic explanations, for instance, aim to produce observer-independent *facts* about *how the model works*, assuming that a neural network's complex behaviour can be reduced to a simpler, human-legible algorithm implemented by a relatively sparse subset of the network [88]. However, this belief might be grounded on fallacious premises, which mechanistic interpretability researchers themselves recognise [106], acknowledging that "many of the most strongly-connected features are hard to interpret" [3], that "more often than not [...] a sparse set of latents that encode some useful concept of interest do not exist [sic]. It is unclear what causes this problem. One hypothesis is that the concept we want isn't how the model 'thinks' about the concept" [106] and that their methods struggle telling apart *faithful* from *plausible* explanations [106].⁵

4.1.3 Ante-hoc methods: the Death of the Model. The onto-epistemological implications of *ante-hoc* methods deserve special attention. The simplified/reduced/pruned models are not merely *surrogates*, as they are often framed, but, in Baudrillard's terms, *simulacra* [8]: copies without an original. They are plausible, human-friendly projections obtained through layered reductionist operations (abstraction, decomposition, pruning, selection). These simulacra pretend to represent the model's logic, yet their grounding is tenuous. These methods offer the comfort of recognising familiar meaning in an alien, possibly unexplainable system, and thus function as epistemological palliatives, structured ways of soothing the desire

⁵This limitation aligns with Carboni and colleagues' comment of *fauxomation*, the false sense of explanation staged by XAI techniques, which "often offload ambiguity back to human users" [22].

for clarity in the face of ungraspable inhuman complexity. However, like the ethnologist in Baurdillard's story, whose observations cause their object (indigenous people) to vanish [8], the AI model also "disintegrates immediately upon contact" because in "order for ethnology to live, its object must die; by dying (it) defies the science that wants to grasp it". In the case of ante-hoc methods, the act of forcing the model into a simplified, interpretable form erases the original system, leaving behind only a projection supposedly tailored to human understanding. The model vanishes, and what remains interpretable is not the model's actual logic but the residue of the reductionist operations that made it legible. Some mechanistic interpretability researchers are aware of what they call *mechanistic faithfulness*; i.e., the epistemic gap between the replacement and the original model: "we cannot guarantee that the replacement model has learned the same mechanisms" [3]. These methods are indeed based on "convenient assumptions", like "that only active features are involved in the mechanism underlying a model's responses" [3] or on self-identified "speculative claims and assumptions about neural networks" like that "features can be rigorously studied and understood" [93].

4.1.4 Human-machine Incommensurability. The various theoretical issues we identified with the XAI methods discussed in this sections might not be easily solved with a technical fix as they might instead stem from the fundamentally different grounds that machines and humans have, which are not directly accessible to the other, what Fazi [40] calls *incommensurability*. Fazi explains that human and machine abstractions are incommensurable: they "cannot be measured against each other or compared by a common standard". She contends that XAI aims to bring "what is beyond human knowledge back into the domain of human cognitive representation", but warns that "opening the black box" might reveal "nothing to translate or to render precisely because the possibility of human representation never existed in the first place". Fazi argues that the way in which AI learns features is indeed entirely and exclusively computational and, thus, we cannot have full epistemological access to it. Alvarado echoes this sentiment, suggesting that humans cannot have epistemic access to features that are epistemically relevant for machines, and thus "explainability is simply not possible" [2]. Similarly, Burrell identified the flaws in any attempt to impose "a process of human interpretive reasoning on a mathematical process of statistical optimization" [21].

These numerous onto-epistemological issues and tenuous assumptions motivate us to look beyond reflectivity and refractivity and to assess whether a diffractive⁶ optic offers a more philosophically robust reading of XAI practices.

4.2 A Diffractive Reading of XAI within Agential Realism

In agential realism, knowledge does not pre-exist the practices and the beings that produce it. Similarly, in an agential realism reading of XAI, interpretation is not a unidirectional transfer of pre-existing

⁶As explained in Section 2.3, diffraction, as introduced by Barad, is not simply another optic alongside refraction and reflection but the guiding practice through which her broader framework of agential realism is articulated. Thus, when we speak of an agential-realist onto-epistemology of XAI, we mobilise both the optic of diffraction and the agential realist commitments it entails: intra-action, apparatuses, agential cuts, and the performativity of knowledge practices.

meaning, but a relational co-production of interpretable *phenomena*. These phenomena emerge through intra-actions between entanglements of human and non-human agencies and actants. This idea resonates with what Ehsan and colleagues [32] have previously argued about an explanation not being merely a function of the AI system, but a shared meaning-making process situated within social practices. The ability to explain does not reside solely “in the box,” but emerges from who is around the box and what tools they have. On this view, interpretations are *emergent* rather than immanent.

4.2.1 Performing Emergent Explanations. Under agential realism, XAI is a *material-discursive* performance (Table 1c): it involves technical tools (e.g., methods, approximations, visualisation tools) as well as the narratives and contexts that constitute what is made intelligible and what is excluded from intelligibility. A *heatmap* explanation for an image classifier, for instance, is a material artefact (pixels highlighted on a screen) that only becomes meaningful through discourse: “the model focuses on this region because it has learned to associate these pixels with the class *cat*”. The discursive aspect involves attending to such explanations as emerging through different entangled apparatuses (the tools and methods in their situated configuration, Eq. 5).

Each interpretive apparatus enacts an agential cut (χ in Eq. 6). For example, in *interactive debugging* and *explanatory debugging*, this cut is realised through explicit interaction loops: users inspect outputs and adjust inputs or parameters, so that interpretations arise from the mutual dependence between user actions, model behaviour, and interface representations. *Counterfactual perturbations* enact cuts in which human interpreters are folded into the explanatory practice as alternative inputs are reasoned with by them. Similarly, *contrastive explanations* are situated: the very question of “why this rather than that” positions the human interpreter within the explanation, foregrounding what makes one explanation possible against possible alternatives and thus aligning the model’s distinctions with contrasts that the interpreter can perceive and evaluate. Although often presented as a method of locating linguistic knowledge within a model, *probing* can be seen as an agential cut that selects certain features (syntax, roles, entities) as salient while excluding other possible relationalities, and constructs significance by mapping researcher-defined linguistic categories onto the model’s internal representations.

4.2.2 Analysing Diffractive Signatures. A diffractive stance stages these readings together and treats interpretation as *reading patterns of difference* through one another (Eqs. 6 and 7). Brought into deliberate juxtaposition, explanation traces interfere with each other and produce possible interpretations. This juxtaposition can be read through opposition theory, which holds that meaning arises from contrasts rather than from intrinsic properties [100]: what a feature or explanation “is” becomes intelligible through the differences it makes across different interpretations. Divergences and alignments co-produce an *interference pattern* of intelligibility, where some interpretations are reinforced while others are cancelled and candidate explanations gain credibility where different methods repeatedly converge. The resulting *diffractive signature* is contained in a patterned interrelation of situated explanations rather than in any single one. What counts as an explanation (Eq.7) therefore

emerges through the configuration of model, method, interface, and observer. For instance, a *saliency map* read through a *perturbation analysis* directs attention to features that neither method would have singled out independently.

4.2.3 Co-constituting Agencies via Boundary-making Performances. Agential realism expands the epistemological focus of XAI to include all the agencies constituted in the interpretive performance, not just the model and the explanation. As suggested by Orlikowski and Scott, “if we take [inseparability] seriously the primary unit for research is not independent objects with inherent boundaries and properties but phenomena materially enacted in practice” [94]. Thus, the *explanandum*, the explainer, and the explaining apparatus do not precede but are constituted through the act of interpretation. The explainer does not look at the system from outside of it but is entangled in it; their role and capacity to explain are also enacted within this practice, through which they become a different kind of knowing subject (Eq. 8). The apparatus itself is also constituted: there is no stable *edge detector neuron*, but only an observed activation that appears related to edges through prior knowledge and technical conventions.

4.2.4 Opening to Multiple and Uncertain Interpretations. Diffractive XAI not only legitimises but demands emerging interpretations to be *one among others*, not *the* interpretation as different interpretations can and do emerge from different entanglements. Confirmations for this idea come from Nicenboim and colleagues [90], who explained that individual explanations cannot fit all contexts and thus there is a need to design for multiple interpretations to be produced, and from Meloux and colleagues’ empirical finding that multiple and even incompatible explanations can be generated with a single method [88]. These entanglements are akin to Bohr’s experimental arrangements: each reveals one facet while obscuring others. Plurality and situatedness of interpretations stand in sharp contrast to traditional positions in XAI, which tend to treat ambiguity and multiplicity as problems to be solved in pursuit of a singular, sanitised account of model behaviour. Mechanistic interpretability research indeed assumes *explanatory unicity*; the hypothesis that a single explanation exists for a given phenomenon [88]. Interpretability also cannot fully eliminate uncertainty because, a feature of relational sense-making and not an obstacle to it. As commented by Carboni and colleagues, “opacity and uncertainty are not deficits to be eliminated but effects of relational configurations of humans, data, models, and institutional practices” [22].

5 Ethical Implications

The mainstream XAI research paradigm downplays researchers’ ethical responsibility, as their role to recover or reveal a model’s pre-existing internal structure seems ethically inconsequential. From a Baradian standpoint, however, the onto-epistemological practices through which XAI itself takes shape have three distinct implications.

First, agential cuts that constitute explanatory performances are *acts of responsibility* [99], as they participate in making some interpretations possible while foreclosing others [22]. Agential cuts are moments in which specific entanglements of models, data, funders,

researchers, and institutions are stabilised so that some interpretations come to matter, while others recede. Responsibility for these cuts must be considered and their consequences made visible [41]. Following Fuchsberger and Frauenberger [42], in more-than-human assemblages, responsibility is not static and distributed, but is performed, relational, fluid, and collectively produced; it is *done*. Doing responsibility in XAI means rejecting the idea that responsibilities can be cleanly distributed to individual components of the explanatory entanglements and revealing the imaginary underlying specific entanglements, the funding flows they follow, and the politics they are aligned with, and the power asymmetries embedded within entanglements. An analysis of entanglements that include mechanistic interpretability, for instance, reveals that the field is currently dominated by a small number of research institutions that set the political agenda, which sees interpretability as a goal to be pursued at all costs, despite researchers from these very same institutions warning that progress in interpretability research has been modest and has not yet demonstrably made AI systems safer [106]. Responsibilities indeed begin in the milieu of innovation that make particular technological futures imaginable and viable [42]. In the case of mechanistic interpretability, for instance, such a milieu can be partially reconstructed by following the funding that flows into this research area, the affiliations of its researchers, who mostly work for companies associated with the TESCREAL bundle [45], and their policy-facing ambitions to “anticipate dangerous capabilities” and to “translate technical progress into levers for governance” [106].

Second, XAI researchers are accountable for identifying and correcting the patterns of exclusion caused by their explanatory entanglements. Giraud’s *ethics of exclusion* critiques entanglement theories that fail to identify who is systematically left out of entanglements [49]. Entanglements do not ensure ethical adequacy so they must be interrogated [57] to surface which explanation methods and metrics are privileged, which kinds of expertise and reporting venues define what counts as a valid explanation, and which individuals and communities affected by model decisions are marginalised. When these exclusions persist, transparency becomes a hollow demand: visibility does not yield accountability if those affected lack the power to make claims or seek remedies, as critiques of the transparency ideal in algorithmic governance have shown [5]. For communities that already experience epistemic exclusion (i.e., by being prevented from participating in collective meaning-making and epistemic processes [59]), the risk is that entanglements might exacerbate these issues. Framing exclusion as a matter of accountability therefore re-politicises interpretability by asking who is entitled to receive and contest explanations, who must answer for the consequences of model use, and how these obligations are distributed [110].

Third, agential realism offers new critical tools to challenge the vernacular understanding of transparency. Reducing transparency to a matter of black-box opacity naively simplifies questions of accountability [5, 110], as visibility alone cannot reveal the sociotechnical context that conditions model behaviour. From an agential realist perspective, transparency is always partial and situated, and is mediated by the specific entanglements that make *seeing* possible. For instance, documentation and analytic practices such as datasheets for datasets, model cards, and genealogies of data work

show that making systems legible is not a simple matter of disclosure but of curating accounts of a system’s life cycle [28, 80, 101]. These interventions involve situated decisions about which histories, assumptions, and data lineages to foreground and which to omit, thereby stabilising particular accounts of where data come from and who is implicated in their production [28, 85]. These documentation practices are material-discursive configurations that actively participate in the production of accountability.

A diffractive approach therefore reframes (a) *responsibility* as an ongoing, collective practice of making, tracing, and contesting agential cuts within explanatory entanglements, and (b) *accountability* as the situated reconfiguration of who can see, contest, and demand redress for model behaviour within the apparatus of models, explanations, and institutions.

6 Diffractive Design for Emergent Interpretations

While developing a comprehensive set of concrete design prescriptions is beyond the scope of this paper, this section outlines some directions to provide a starting point for designing diffractive interfaces that support emergent interpretations. These directions are informed by several works in HCI, and, in particular, by the aforementioned work on diffractive interfaces [47] and by studies on emergence, which is central in our proposed onto-epistemology. Gaver and colleagues articulated how, in design research, the objectives, procedures, and outcomes often do not follow the execution of a fixed plan but *emerge* through making [43]. Building on this view, more recent HCI work focused on understanding how AI systems are configured and emerge through ongoing relations with artefacts and environments. For example, Nicenboim and colleagues used conversational agents and more-than-human prototypes to stage open-ended situations in which partial, situated, and often mistaken understandings of AI can evolve over time [91, 92]. In parallel, Ghajargar and Bardzell’s “graspable AI” concept cards act as tangible prompts through which designers’ ideas about AI can iteratively emerge [46].

In the remainder of this section, after presenting each design direction, we illustrate it with a practical example: a persona interacting with a speculative interface [31] designed to explain the outputs of a generative text-to-music (TTM) system. The interface allows users to generate songs using the AI model and to query the outputs to develop specific interpretations. In line with our framework, the users can define the configurations of the explanatory entanglements and the generations that they create (e.g., XAI methods, training dataset, explainer’s background, model version, prompt). We chose this case study because it draws on our disciplinary expertise and because generative music offers an especially fertile setting for investigating diffractive XAI. Assessing the constitutions of explanatory entanglements in TTM is indeed particularly important because: i) the training data used in these systems are indeed often scraped without consent and encode a narrow, biased canon of music [17]; ii) the space of allowed outputs is reduced by the epistemic decisions behind these models [83]; and iii) their outputs risk actively displacing human labour [82].

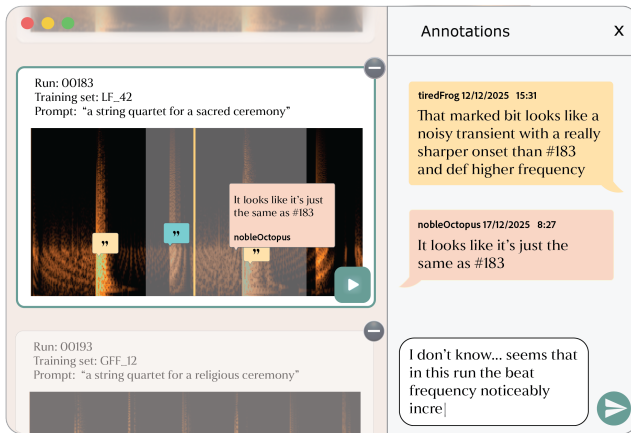


Figure 2: A speculative interface designed to interpret text-to-music generations. It displays multiple situated entanglements annotated by different users.

6.1 Foreground Multiple Interpretations

Supported by findings from related work [88, 90], in Section 4.2.4, we explained how diffractive XAI demands multiple interpretations. Diffractive XAI interfaces should thus refrain from visualising and settling on one correct interpretation and should instead show patterns of difference across several explanatory performances. Diffractive XAI interfaces should not only enable users to select different entanglements and agential cuts, but also stage situated outputs against one another so that they can be experienced in contrast, echoing our earlier point that meaning emerges through comparison rather than from isolated, intrinsic properties.

A model developer aims to understand how the TTM model has internalised the concept of “sacred” in music. They initially fix a specific entanglement and generate a song from a certain prompt (e.g., “a string quartet for a sacred ceremony”). Then, the model produce counterfactuals by minimally varying components of the entanglement like the prompt (e.g., replacing “sacred” with “religious”) or the subset of the training catalogue (e.g., muting tracks from a given region or label). Different users then annotate how sacredness transforms across these versions (Fig. 2). An explainer then analyses the interference pattern between these explanatory performances. The pattern might reveal that sacral and religious co-activate with specific classes of instruments (e.g., organ, choir) that only appear when the training catalogue contains songs from a particular region. The material-discursive performance can thus reveal, for instance, how specific catalogues perpetuate colonial notions of sacral music.

6.2 Foreground Situatedness and Repeated Engagement

In a diffractive view, situatedness and emergence mutually determine interpretations. Explanations arise from partial, situated entanglements and thus cannot be treated in isolation. Rather, they unfold through repeated engagements in which agential cuts shift

and new relations become salient. Building on work on diffractive interfaces [47] and practice-based accounts of emergence [43], diffractive XAI interfaces should foreground the situatedness of the entanglements that gave rise to explanations, render agential cuts visible, and allow users to move between them. They should then support revisiting prior explanations and navigating patterns of differences from successive interactions to perform new interpretations. On this view, diffractive XAI interfaces are spaces where explanations continually emerge and re-emerge through changing and evolving relations among the components of the explanatory entanglement.

An IP lawyer aims to identify which training song played the greatest role in a specific generation (attribution problem). They perform different generations using the same seed⁷ but varying aspects of the entanglements. Initially, they explore different attribution methods (e.g., influence functions, and concept-based probe for “compositional style”) and then they select a different subset of the training dataset, and a combination of these. The user can also analyse candidate songs both visually (by analysing a spectrogram or other visualisation graphs) and acoustically. Each generation-output-explanation cycle is a revisitable interpretive performance related to the specific situatedness of the entanglement. By navigating traces of successive runs as they accumulate over time, and by allowing users to toggle between different entanglements, a candidate explanation emerges.

6.3 Foreground Uncertainty and Ambiguity

Recent work invited to stage misunderstandings and breakdowns in human-AI interaction as occasions for reflection [91, 92]. This suggestion resonates with long-standing HCI work that treats ambiguity and open meaning as resources for design and interpretation [14, 44, 84, 96, 105, 114] and with the design concept of *seamful design* [24], which foregrounds gaps, seams, and incongruities so that multiple readings remain available. Brought to XAI, this principle invites designers to treat uncertainty, ambiguity, non-knowledge, and disagreement as materials for interpretation that should be foregrounded rather than sanitised.

A musician aims to identify and make sense of disagreements between various explanations of the outputs of a TTM model. Specifically, they implement probing heads to map musical concepts they define (e.g., melody, texture, rhythm) onto the model’s internal representations. Rather than averaging these scores to present a single interpretation, the interface deliberately seeks out passages where these probes disagree. For instance, it identifies disagreement where one probe confidently labels a strings sequence as the main melody, while another strongly classifies it as accompaniment, and a third indicates high activation for background texture. The interface foregrounds this conflict and presents these competing interpretations as opportunities for musical inquiry, which might

⁷Generative diffusion models often exhibit a “seed consistency” property: with a fixed architecture and different but related training data, models can produce very similar outputs when given the same prompt and random seed [60].

result in the musician redefining the existing concept labels or inventing entirely new ones. In doing so, musicians are not only active participants in the material-discursive practice of interpretation but become a “different kind of knowing subject” (Section 4.2.3), as their musical expertise changes after participating in the explanatory performance.

7 Conclusions

In this paper, we challenged the assumptions of mainstream XAI research about the existence of hidden *explananda* within AI models that can be recovered by an observer who is external to the interpretative process. Inspired by Baradian’s agential realism, we proposed an alternative onto-epistemology that reframes interpretability as a performed phenomenon that allows possible, but not definitive, explanations to emerge from an entanglement of models, observers, tools, and contexts. Understanding interpretability as a material-discursive performance co-constituted by situated entanglements of human and non-human actors carries not only onto-epistemic but also ethical stakes. Configuring entanglements is an act of responsibility as each entanglement foregrounds some perspectives while forecloses others. XAI researchers are therefore accountable for these configurations as constitutive conditions that must be made visible, acknowledged, and politicised. XAI research, we argue, might benefit from embracing this onto-epistemological shift: not extracting truths from models, but practicing explanations that remain attentive to difference, interference, and accountability. As for the limitations of this study, our framework is primarily conceptual, and we do not present a full empirical case study of diffractive XAI in use. The design directions we outlined have not been developed into operational design guidelines nor evaluated in fully developed scenarios. Future work should therefore translate these directions into concrete design patterns and guidelines.

8 Ethical Statement

In this paper, we cited self-published technical works only insofar as they provided insights into their methods and thinking, treating them strictly as secondary data sources. Other non-peer-reviewed literature has not been considered or referenced in order to uphold the standards of academic rigour.

References

- [1] Suzana Alpsancar, Heike M. Buhl, Tobias Matzner, and Ingrid Scharlau. 2024. Explanation needs and ethical demands: unpacking the instrumental value of XAI. *AI and Ethics* (Dec. 2024). doi:10.1007/s43681-024-00622-3
- [2] Ramón Alvarado. 2023. AI as an Epistemic Technology. *Science and Engineering Ethics* 29, 5 (Oct. 2023), 32. doi:10.1007/s11948-023-00451-3
- [3] Emmanuel Ameisen, Jack Lindsey, Adam Pearce, Wes Gurnee, Nicholas L Turner, Brian Chen, Craig Citro, David Abrahams, Shan Carter, Basil Hosmer, and others. 2025. Circuit tracing: Revealing computational graphs in language models. *Transformer Circuits Thread* 6 (2025).
- [4] Saleema Amershi, Max Chickering, Steven Drucker, Bongshin Lee, Patrice Simard, and Jina Suh. 2015. ModelTracker: Redesigning Performance Analysis Tools for Machine Learning. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI 2015)*. ACM.
- [5] Mike Ananny and Kate Crawford. 2018. Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media & Society* 20, 3 (March 2018), 973–989. doi:10.1177/1461444816676645
- [6] Karen Barad. 2003. Posthumanist Performativity: Toward an Understanding of How Matter Comes to Matter. *Signs: Journal of Women in Culture and Society* 28, 3 (March 2003), 801–831. doi:10.1086/345321
- [7] Karen Barad. 2007. Meeting the Universe Halfway: Quantum Physics and the Entanglement of Matter and Meaning. (2007), 542.
- [8] Jean Baudrillard. 1994. *Simulacra and simulation*. University of Michigan Press, Ann Arbor.
- [9] Nora Belrose, Zach Furman, Logan Smith, Danny Halawi, Igor Ostrovsky, Lev McKinney, Stella Biderman, and Jacob Steinhardt. 2023. Eliciting Latent Predictions from Transformers with the Tuned Lens. doi:10.48550/arXiv.2303.08112 arXiv:2303.08112 [cs].
- [10] Astrid Bertrand, Tiphaine Viard, Rafik Belloum, James R. Eagan, and Winston Maxwell. 2023. On Selective, Mutable and Dialogic XAI: a Review of What Users Say about Different Types of Interactive Explanations. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3544548.3581314 event-place: Hamburg, Germany.
- [11] Aditya Bhattacharya, Simone Stumpf, Robin De Croon, and Katrien Verbert. 2025. Explanatory Debiasing: Involving Domain Experts in the Data Generation Process to Mitigate Representation Bias in AI Systems. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3706598.3713497
- [12] Steven Bills, Nick Cammarata, Dan Mossing, and Leo Gao. 2023. Language models can explain neurons in language models. <https://openaipublic.blob.core.windows.net/neuron-explainer/paper/index.html>
- [13] Charlotte Bird. 2025. Diffraction, Creativity and AI: Towards New Methods for Design Research. (2025).
- [14] Kirsten Boehner and Jeffrey T. Hancock. 2006. Advancing ambiguity. In *Proceedings of the SIGCHI conference on Human Factors in computing systems - CHI '06*. ACM Press, Montreal, Quebec, Canada, 103. doi:10.1145/1124772.1124789
- [15] Niels Bohr. 1928. The Quantum Postulate and the Recent Development of Atomic Theory. *Nature* 121 (1928), 580–590.
- [16] Federico Bomba, María Menéndez-Blanco, Paolo Grigis, Michele Cremaschi, and Antonella De Angeli. 2024. The Choreographer-Performer Continuum: A Diffraction Tool to Illuminate Authorship in More Than Human Co-Performances. *ACM Transactions on Computer-Human Interaction* 31, 6 (Dec. 2024), 1–23. doi:10.1145/3689040
- [17] Georgina Born. 2020. Diversifying MIR: Knowledge and Real-World Challenges, and New Interdisciplinary Futures. *Transactions of the International Society for Music Information Retrieval* 3, 1 (Oct. 2020), 193–204. doi:10.5334/tismir.58
- [18] Trenton Bricken, Adly Templeton, Joshua Batson, Brian Chen, Adam Jermyn, Tom Conerly, Nick Turner, Cem Anil, Carson Denison, Amanda Askell, and others. 2023. Towards monosemanticity: Decomposing language models with dictionary learning. *Transformer Circuits Thread* 2 (2023).
- [19] Jude Browne, Stephen Cave, Eleanor Drage, and Kerry McInerney. 2023. *Feminist AI: Critical Perspectives on Algorithms, Data, and Intelligent Machines* (1 ed.). Oxford University Press/Oxford. doi:10.1093/oso/9780192889898.001.0001
- [20] Nick Bryan-Kinns, Shuoyang Zheng, Francisco Castro, Makayla Lewis, Jia-Rey Chang, Gabriel Vigliensoni, Terence Broad, Michael Paul Clemens, and Elizabeth Wilson. 2025. XAIxArts Manifesto: Explainable AI for the Arts. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–8. doi:10.1145/3706599.3716227
- [21] Jenna Burrell. 2016. How the machine ‘thinks’: Understanding opacity in machine learning algorithms. *Big Data & Society* 3, 1 (June 2016), 2053951715622512. doi:10.1177/2053951715622512
- [22] Chiara Carboni, Rik Wehrens, Romke Van Der Veen, and Antoinette De Bont. 2023. Eye for an AI: More-than-seeing, fautoimation, and the enactment of uncertain data in digital pathology. *Social Studies of Science* 53, 5 (Oct. 2023), 712–737. doi:10.1177/03063127231167589
- [23] Dubravka Cecez-Kecmanovic. 2016. From Substantialist to Process Metaphysics – Exploring Shifts in IS Research. In *Beyond Interpretivism? New Encounters with Technology and Organization*, Lucas Introna, Donncha Kavanagh, Séamas Kelly, Wanda Orlikowski, and Susan Scott (Eds.). Vol. 489. Springer International Publishing, Cham, 35–57. doi:10.1007/978-3-319-49733-4_3 Series Title: IFIP Advances in Information and Communication Technology.
- [24] Matthew Chalmers and Areti Galani. 2004. Seamlful interweaving: heterogeneity in the theory and design of interactive systems. In *Proceedings of the 2004 conference on Designing interactive systems processes, practices, methods, and techniques - DIS '04*. ACM Press, Cambridge, MA, USA, 243. doi:10.1145/1013115.1013149
- [25] Chaofan Chen, Oscar Li, Daniel Tao, Alina Barnett, Cynthia Rudin, and Jonathan K Su. 2019. This Looks Like That: Deep Learning for Interpretable Image Recognition. In *Advances in Neural Information Processing Systems*, Vol. 32. Curran Associates, Inc.
- [26] Marianela Ciolfi Felice, Ivana Feldfeber, Carolina Glasserman Apicella, Yasmín Belén Quiroga, Julián Ansaldo, Luciano Lapenna, Santiago Bezchinsky, Raul Barriga Rubio, and Mailén García. 2025. Doing the Feminist Work in AI: Reflections from an AI Project in Latin America. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–18. doi:10.1145/3706598.3713681

- [27] Regina de Brito Duarte, Mónica Costa Abreu, Joana Campos, and Ana Paiva. 2025. The Amplifying Effect of Explainability in AI-assisted Decision-making in Groups. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3706598.3713534
- [28] Emily Denton, Alex Hanna, Razvan Amironesei, Andrew Smart, and Hilary Nicole. 2021. On the genealogy of machine learning datasets: A critical history of ImageNet. *Big Data & Society* 8, 2 (July 2021), 205395172110359. doi:10.1177/20539517211035955
- [29] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *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)*. 4171–4186.
- [30] Eleanor Drage, Kerry McInerney, and Jude Browne. 2024. Engineers on responsibility: feminist approaches to who's responsible for ethical AI. *Ethics and Information Technology* 26, 1 (March 2024), 4. doi:10.1007/s10676-023-09739-1
- [31] Anthony Dunne and Fiona Raby. 2013. *Speculative everything: design, fiction, and social dreaming*. The MIT Press, Cambridge, Massachusetts ; London.
- [32] Upol Ehsan, Q. Vera Liao, Michael Muller, Mark O. Riedl, and Justin D. Weisz. 2021. Expanding Explainability: Towards Social Transparency in AI systems. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3411764.3445188
- [33] Upol Ehsan, Samir Passi, Q. Vera Liao, Larry Chan, I-Hsiang Lee, Michael Muller, and Mark O Riedl. 2024. 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 (CHI '24)*. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3613904.3642474 event-place: Honolulu, HI, USA.
- [34] Upol Ehsan and Mark O. Riedl. 2024. Social construction of XAI: Do we need one definition to rule them all? *Patterns* 5, 2 (Feb. 2024). doi:10.1016/j.patter.2024.100926 Publisher: Elsevier.
- [35] Upol Ehsan, Koustuv Saha, Mummun De Choudhury, and Mark O Riedl. 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 (2023), 1–32. Publisher: ACM New York, NY, USA.
- [36] Upol Ehsan, Elizabeth A Watkins, Philipp Wintersberger, Carina Manger, Nina Hubig, Saiph Savage, Justin D. Weisz, and Andreas Riemer. 2025. New Frontiers of Human-centered Explainable AI (HCXAI): Participatory Civic AI, Benchmarking LLMs, XAI Hallucinations, and Responsible AI Audits. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25)*. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3706599.3706713
- [37] Will Epperson, Gagan Bansal, Victor C Dibia, Adam Fournay, Jack Gerrits, Erkang Zhu, and Saleema Amershi. 2025. Interactive debugging and steering of multi-agent ai systems. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [38] Adrian Erasmus, Tyler D. P. Brunet, and Eyal Fisher. 2021. What is Interpretability? *Philosophy & Technology* 34, 4 (Dec. 2021), 833–862. doi:10.1007/s13347-020-00435-2
- [39] Eva Eriksson, Daisy Yoo, Tilde Bekker, and Elisabet M. Nilsson. 2024. More-than-Human Perspectives in Human-Computer Interaction Research: A Scoping Review. In *Nordic Conference on Human-Computer Interaction*. ACM, Uppsala Sweden, 1–18. doi:10.1145/3679318.3685408
- [40] M. Beatrice Fazi. 2021. Beyond Human: Deep Learning, Explainability and Representation. *Theory, Culture & Society* 38, 7-8 (Dec. 2021), 55–77. doi:10.1177/0263276420966386
- [41] Christopher Frauenberger. 2019. Entanglement HCI The Next Wave? *ACM Transactions on Computer-Human Interaction* 27, 1 (2019), 1–27. doi:10.1145/3364998
- [42] Verena Fuchsberger and Christopher Frauenberger. 2025. Doing responsibilities in entangled worlds. *Human-Computer Interaction* 40, 1-4 (July 2025), 65–88. doi:10.1080/07370024.2023.2269934
- [43] William Gaver, Peter Gall Krogh, Andy Boucher, and David Chatting. 2022. Emergence as a Feature of Practice-based Design Research. In *Designing Interactive Systems Conference*. ACM, Virtual Event Australia, 517–526. doi:10.1145/3532106.3533524
- [44] William W Gaver, Jacob Beaver, and Steve Benford. 2003. Ambiguity as a Resource for Design. *NEW HORIZONS* 5 (2003), 8.
- [45] Timnit Gebru and Émile P. Torres. 2024. The TESCREAL bundle: Eugenics and the promise of utopia through artificial general intelligence. *First Monday* (April 2024). doi:10.5210/fm.v29i4.13636
- [46] Maliheh Ghajargar and Jeffrey Bardzell. 2022. Making AI Understandable by Making it Tangible: Exploring the Design Space with Ten Concept Cards. In *Proceedings of the 34th Australian Conference on Human-Computer Interaction*. ACM, Canberra ACT Australia, 74–80. doi:10.1145/3572921.3572942
- [47] Elisa Giaccardi, Seowoo Nam, and Iohanna Nicenboim. 2025. Diffractive Interfaces: Facilitating Agential Cuts in Forest Data Across More-than-human Scales. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference*. ACM, Madeira Portugal, 135–147. doi:10.1145/3715336.3735404
- [48] Leilani H Gilpin, David Bau, Ben Z Yuan, Ayesha Bajwa, Michael Specter, and Lalana Kagal. 2018. Explaining explanations: An overview of interpretability of machine learning. In *2018 IEEE 5th International Conference on data science and advanced analytics (DSAA)*. IEEE, 80–89.
- [49] Eva Haifa Giraud. 2019. *What comes after entanglement? activism, anthropocentrism, and an ethics of exclusion*. Duke University Press, Durham.
- [50] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and Harnessing Adversarial Examples.
- [51] Leif Hancock-Li and I. Elizabeth Kumar. 2021. Epistemic values in feature importance methods: Lessons from feminist epistemology. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Virtual Event Canada, 817–826. doi:10.1145/3442188.3445943
- [52] Donna Haraway. 1992. The promises of monsters: A regenerative politics for inappropriate/d others. In *Cultural studies*. Routledge, 295–337.
- [53] Donna Jeanne Haraway. 1988. Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective. (1988).
- [54] Cher M Hill. 2018. More-than-reflective practice: Becoming a diffractive practitioner. (2018).
- [55] Patricia Hill Collins. 2000. *Black feminist thought: knowledge, consciousness, and the politics of empowerment* (revision tenth anniversary edition ed.). Routledge, New York. doi:10.4324/9780203900055
- [56] Fred Hohman, Andrew Head, Rich Caruana, Robert DeLine, and Steven M Drucker. 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.
- [57] Gregory Hollin, Isla Forsyth, Eva Giraud, and Tracey Potts. 2017. (Dis)entangling Barad: Materialisms and ethics. *Social Studies of Science* 47, 6 (Dec. 2017), 918–941. doi:10.1177/0306312717728344
- [58] The White House. 2025. Winning the Race: America's AI Action Plan. (2025).
- [59] Linus Ta-Lun Huang, Hsiang-Yun Chen, Ying-Tung Lin, Tsung-Ren Huang, and Tzu-Wei Hung. 2022. Ameliorating Algorithmic Bias, or Why Explainable AI Needs Feminist Philosophy. *Feminist Philosophy Quarterly* 8, 3/4 (Dec. 2022). doi:10.5206/fpq/2022.3/4.14347
- [60] Zahra Kadhodaie, Florentin Guth, Eero P. Simoncelli, and Stéphane Mallat. 2024. Generalization in diffusion models arises from geometry-adaptive harmonic representations. doi:10.48550/arXiv.2310.02557 arXiv:2310.02557 [cs].
- [61] Been Kim, Martin Wattenberg, Justin Gilmer, Carrie Cai, James Wexler, Fernanda Viegas, and Rory Sayres. 2018. Interpretability Beyond Feature Attribution: Quantitative Testing with Concept Activation Vectors (TCAV). doi:10.48550/arXiv.1711.11279 arXiv:1711.11279 [stat].
- [62] Jinkyu Kim and John Canny. 2017. Interpretable Learning for Self-Driving Cars by Visualizing Causal Attention. doi:10.48550/arXiv.1703.10631 arXiv:1703.10631 [cs].
- [63] Sunnie S. Y. Kim, Elizabeth Anne Watkins, Olga Russakovsky, Ruth Fong, and Andrés Monroy-Hernández. 2023. "Help Me Help the AI": Understanding How Explainability Can Support Human-AI Interaction. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3544548.3581001 event-place: Hamburg, Germany.
- [64] Goda Klumbyté. 2023. Towards Feminist Intersectional XAI: From Explainability to Response- Ability. (2023).
- [65] Goda Klumbyté, Hannah Piehl, and Claude Draude. 2023. Explaining the ghosts: Feminist intersectional XAI and cartography as methods to account for invisible labour. (2023).
- [66] Pang Wei Koh and Percy Liang. 2020. Understanding Black-box Predictions via Influence Functions. doi:10.48550/arXiv.1703.04730 arXiv:1703.04730 [stat].
- [67] Josua Krause, Adam Perer, and Kenney Ng. 2016. Interacting with Predictions: Visual Inspection of Black-box Machine Learning Models. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. Association for Computing Machinery, New York, NY, USA, 5686–5697. doi:10.1145/2858036.2858529
- [68] Maya Krishnan. 2020. Against Interpretability: a Critical Examination of the Interpretability Problem in Machine Learning. *Philosophy & Technology* 33, 3 (Sept. 2020), 487–502. doi:10.1007/s13347-019-00372-9
- [69] Todd Kulesza, Margaret Burnett, Weng-Keen Wong, and Simone Stumpf. 2015. Principles of Explanatory Debugging to Personalize Interactive Machine Learning. In *Proceedings of the 20th International Conference on Intelligent User Interfaces (IUI '15)*. Association for Computing Machinery, New York, NY, USA, 126–137. doi:10.1145/2678025.2701399 event-place: Atlanta, Georgia, USA.
- [70] Lena Kästner and Barnaby Crook. 2024. Explaining AI through mechanistic interpretability. *European Journal for Philosophy of Science* 14, 4 (Dec. 2024), 52. doi:10.1007/s13194-024-00614-4
- [71] Bruno Latour. 1987. *Science in action: how to follow scientists and engineers through society*. Harvard University Press, Cambridge, Mass.

- [72] Susanne Lettow. 2017. Turning the turn: New materialism, historical materialism and critical theory. *Thesis Eleven* 140, 1 (June 2017), 106–121. doi:10.1177/0725513616683853
- [73] Q. Vera Liao, Hariharan Subramonyam, Jennifer Wang, and Jennifer Wortman Vaughan. 2023. Designly Understanding: Information Needs for Model Transparency to Support Design Ideation for AI-Powered User Experience. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, Hamburg Germany, 1–21. doi:10.1145/3544548.3580652
- [74] Q. Vera Liao, Mihaela Vorvoreanu, Hari Subramonyam, and Lauren Wilcox. 2024. UX Matters: The Critical Role of UX in Responsible AI. *Interactions* 31, 4 (July 2024), 22–27. doi:10.1145/3665504
- [75] Pantelis Linardatos, Vasilis Papastefanopoulos, and Sotiris Kotsiantis. 2020. Explainable AI: A Review of Machine Learning Interpretability Methods. *Entropy* 23, 1 (Dec. 2020), 18. doi:10.3390/e23010018
- [76] Jack Lindsey, Wes Gurnee, Emmanuel Ameisen, Brian Chen, Adam Pearce, Nicholas L Turner, Craig Citro, David Abrahams, Shan Carter, Basil Hosmer, and others. 2025. On the biology of a large language model. *Transformer Circuits Thread* (2025).
- [77] Scott M Lundberg and Su-In Lee. 2017. A Unified Approach to Interpreting Model Predictions. In *Advances in Neural Information Processing Systems*, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2017/file/8a20a8621978632d76c43dfd28b67767-Paper.pdf
- [78] Lia Mice and Andrew P McPherson. 2022. Super Size Me: Interface Size, Identity and Embodiment in Digital Musical Instrument Design. (2022), 15.
- [79] Tim Miller. 2021. Contrastive Explanation: A Structural-Model Approach. *The Knowledge Engineering Review* 36 (2021), e14. doi:10.1017/S0269888921000102 arXiv:1811.03163 [cs].
- [80] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*. 220–229.
- [81] Grégoire Montavon, Alexander Binder, Sebastian Lapuschkin, Wojciech Samek, and Klaus-Robert Müller. 2019. Layer-wise relevance propagation: an overview. *Explainable AI: interpreting, explaining and visualizing deep learning* (2019), 193–209. Publisher: Springer.
- [82] Fabio Morreale, Elham Bahmanteymouri, Brent Burmester, Andrew Chen, and Michelle Thorp. 2023. The unwitting labourer: extracting humanness in AI training. *AI & SOCIETY* (May 2023). doi:10.1007/s00146-023-01692-3
- [83] Fabio Morreale, Marco A. Martinez-Ramirez, Raul Masu, WeiHsiang Liao, and Yuki Mitsufuji. 2025. Reductive, Exclusionary, Normalising: The Limits of Generative AI Music. *Transactions of the International Society for Music Information Retrieval* (2025). doi:10.5334/tismir.256
- [84] Fabio Morreale, Raul Masu, and Antonella De Angeli. 2019. The Influence of Coauthorship in the Interpretation of Multimodal Interfaces. *Wireless Communications and Mobile Computing* 2019 (2019), 1–12. doi:10.1155/2019/9234812
- [85] Fabio Morreale, Megha Sharma, and I-Chieh Wei. 2023. Data Collection in Music Generation Training Sets: A Critical Analysis.
- [86] Landon Morrison and Andrew McPherson. 2024. Entangling Entanglement: A Diffractive Dialogue on HCI and Musical Interactions. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–17. doi:10.1145/3613904.3642171
- [87] Tom Mudd. 2023. Playing with Feedback: Unpredictability, Immediacy, and Entangled Agency in the No-input Mixing Desk. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, Hamburg Germany, 1–11. doi:10.1145/3544548.3580662
- [88] Maxime Méloux, Silviu Maniu, François Portet, and Maxime Peyrard. 2025. Everything, Everywhere, All at Once: Is Mechanistic Interpretability Identifiable? doi:10.48550/arXiv.2502.20914 arXiv:2502.20914 [cs].
- [89] Romy Müller, Marius Thoß, Julian Ullrich, Steffen Seitz, and Carsten Knoll. 2025. Interpretability is in the Eye of the Beholder: Human Versus Artificial Classification of Image Segments Generated by Humans Versus XAI. *International Journal of Human-Computer Interaction* 41, 4 (Feb. 2025), 2371–2393. doi:10.1080/10447318.2024.2323263
- [90] Johanna Nicenboim, Elisa Giaccardi, and Johan Redström. 2022. From explanations to shared understandings of AI. doi:10.21606/drs.2022.773
- [91] Johanna Nicenboim, Marie Louise Juul Søndergaard, Joseph Lindley, Anuradha Reddy, Yolande Strengers, Johan Redström, and Elisa Giaccardi. 2024. Unmaking with AI: Tactics for Decentering through Design. *ACM Transactions on Computer-Human Interaction* 31, 6 (Dec. 2024), 1–20. doi:10.1145/3685275
- [92] Johanna Nicenboim, Shruthi Venkat, Neva Linn Rustad, Diana Vardanyan, Elisa Giaccardi, and Johan Redström. 2023. Conversation Starters: How Can We Misunderstand AI Better?. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, Hamburg Germany, 1–4. doi:10.1145/3544549.3583914
- [93] Chris Olah, Nick Cammarata, Ludwig Schubert, Gabriel Goh, Michael Petrov, and Shan Carter. 2020. Zoom In: An Introduction to Circuits. *Distill* (2020). doi:10.23915/distill.00024.001
- [94] W. J. Orlikowski. 2010. The sociomateriality of organisational life: considering technology in management research. *Cambridge Journal of Economics* 34, 1 (Jan. 2010), 125–141. doi:10.1093/cje/bep058
- [95] Cecilia Panigutti, Andrea Beretta, Fosca Giannotti, and Dino Pedreschi. 2022. Understanding the impact of explanations on advice-taking: a user study for AI-based clinical Decision Support Systems. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3491102.3502104 event-place: New Orleans, LA, USA.
- [96] Courtney N. Reed, Adan L. Benito, Franco Caspe, and Andrew P. McPherson. 2024. Shifting Ambiguity, Collapsing Indeterminacy: Designing with Data as Baradian Apparatus. *ACM Transactions on Computer-Human Interaction* (Aug. 2024), 3689043. doi:10.1145/3689043
- [97] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. doi:10.48550/arXiv.1602.04938 arXiv:1602.04938 [cs].
- [98] Nicole Robson, Andrew McPherson, and Nick Bryan-Kinns. 2024. Thinking with Sound: Exploring the Experience of Listening to an Ultrasonic Art Installation. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–14. doi:10.1145/3613904.3642616
- [99] Pedro Sanches, Noura Howell, Vasiliki Tsaknaki, Tom Jenkins, and Karey Helms. 2022. Diffraction-in-action: Designly Explorations of Agential Realism Through Lived Data. In *CHI Conference on Human Factors in Computing Systems*. ACM, New Orleans LA USA, 1–18. doi:10.1145/3491102.3502029
- [100] Ferdinand de Saussure. 1916. *Cours de linguistique générale* (Éd. critique, [nachdr. der ausg. 1916] ed.). Payot, Paris.
- [101] Morgan Klaus Scheuerman, Emily Denton, and Alex Hanna. 2021. Do Datasets Have Politics? Disciplinary Values in Computer Vision Dataset Development. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (Oct. 2021), 1–37. doi:10.1145/3476058 arXiv:2108.04308 [cs].
- [102] Johannes Schneider. 2024. Explainable Generative AI (GenXAI): a survey, conceptualization, and research agenda. *Artificial Intelligence Review* 57, 11 (Sept. 2024), 289. doi:10.1007/s10462-024-10916-x
- [103] Hugo Scurto, Baptiste Caramiaux, and Frederic Bevilacqua. 2021. Prototyping Machine Learning Through Diffractive Art Practice. In *Designing Interactive Systems Conference 2021*. ACM, Virtual Event USA, 2013–2025. doi:10.1145/3461778.3462163
- [104] Ramprasaath R. Selvaraju, Abhishek Das, Ramakrishna Vedantam, Michael Cogswell, Devi Parikh, and Dhruv Batra. 2017. Grad-CAM: Why did you say that? doi:10.48550/arXiv.1611.07450 arXiv:1611.07450 [stat].
- [105] Phoebe Sengers and Bill Gaver. 2006. Staying open to interpretation: engaging multiple meanings in design and evaluation. In *Proceedings of the 6th ACM conference on Designing Interactive systems - DIS '06*. ACM Press, University Park, PA, USA, 99. doi:10.1145/1142405.1142422
- [106] Lee Sharkey, Bilal Chughtai, Joshua Batson, Jack Lindsey, Jeff Wu, Lucius Bushnaq, Nicholas Goldowsky-Dill, Stefan Heimersheim, Alejandro Ortega, Joseph Bloom, Stella Biderman, Adria Garriga-Alonso, Arthur Conmy, Neel Nanda, Jessica Rumbelow, Martin Wattenberg, Nandi Schoots, Joseph Miller, Eric J. Michaud, Stephen Casper, Max Tegmark, William Saunders, David Bau, Eric Todd, Atticus Geiger, Mor Geva, Jesse Hoogland, Daniel Murfet, and Tom McGrath. 2025. Open Problems in Mechanistic Interpretability. doi:10.48550/arXiv.2501.16496 arXiv:2501.16496 [cs].
- [107] Linda Tuhiwai Smith. 2012. *Decolonizing methodologies: research and indigenous peoples* (second edition ed.). Zed Books, London. OCLC: 861538625.
- [108] Kacper Sokol and Peter Flach. 2024. Interpretable representations in explainable AI: from theory to practice. *Data Mining and Knowledge Discovery* 38, 5 (Sept. 2024), 3102–3140. doi:10.1007/s10618-024-01010-5 Publisher: Springer Science and Business Media LLC.
- [109] Kacper Sokol and Julia E Vogt. 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*. 1–8.
- [110] Eugenia Stamboliev. 2023. Proposing a Postcritical AI Literacy: Why We Should Worry Less about Algorithmic Transparency and More about Citizen Empowerment. *Media Theory* (2023).
- [111] Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. BERT Rediscovered the Classical NLP Pipeline. doi:10.48550/arXiv.1905.05950 arXiv:1905.05950 [cs].
- [112] Sophie Toupin. 2023. Shaping feminist artificial intelligence. *New Media & Society* (Feb. 2023), 146144482211507. doi:10.1177/14614448221150776
- [113] Sandra Wachter, Brent Mittelstadt, and Chris Russell. 2017. Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *Harv. J.L. & Tech.* 31 (2017), 841. Publisher: HeinOnline.
- [114] Ronald Wakkary. 2021. *Things we could design: for more than human-centered worlds*. The MIT Press, Cambridge, Massachusetts.
- [115] Katharina Weitz, Ruben Schlagowski, Elisabeth André, Maris Männiste, and Ceenu George. 2024. Explaining It Your Way - Findings from a Co-Creative Design Workshop on Designing XAI Applications with AI End-Users from the Public Sector. In *Proceedings of the 2024 CHI Conference on Human Factors in*

- Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3613904.3642563 event-place: Honolulu, HI, USA.
- [116] James Wexler, Mahima Pushkarna, Tolga Bolukbasi, Martin Wattenberg, Fernanda Viégas, and Jimbo Wilson. 2019. The what-if tool: Interactive probing of machine learning models. *IEEE transactions on visualization and computer graphics* 26, 1 (2019), 56–65. Publisher: IEEE.
- [117] Langdon Winner. 1993. Upon Opening the Black Box and Finding It Empty: Social Constructivism and the Philosophy of Technology. *Science, Technology, & Human Values* 18, 3 (July 1993), 362–378. doi:10.1177/016224399301800306
- [118] Jichen Zhu, Pedro Sanches, Vasiliki Tsaknaki, Willem Van Der Maden, and Irene Kaklopoulou. 2025. The Centers and Margins of Modeling Humans in Well-being Technologies. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–16. doi:10.1145/3706598.3713940
- [119] Carsten Østerlund, Kevin Crowston, and Corey Jackson. 2020. Building an Apparatus: Refractive, Reflective, and Diffractive Readings of Trace Data. *Journal of the Association for Information Systems* (2020), 1–22. doi:10.17705/1jais.00590