

Title	Empowering explainable artificial intelligence through case-based reasoning: a comprehensive exploration
Authors	Pradeep, Preeja;Caro-Martínez, Marta;Wijekoon, Anjana
Publication date	2025-09-16
Original Citation	Pradeep, P., Caro-Martínez, M. and Wijekoon, A. (2025) 'Empowering explainable artificial intelligence through case-based reasoning: a comprehensive exploration', IEEE Transactions on Knowledge and Data Engineering, 37(12), pp. 7120–7139. https://doi.org/10.1109/TKDE.2025.3609825
Type of publication	Article (peer-reviewed)
Link to publisher's version	10.1109/TKDE.2025.3609825
Rights	© 2025, the Authors. Published by IEEE. This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/ - https://creativecommons.org/licenses/by/4.0/
Download date	2026-10-11 15:01:46
Item downloaded from	https://hdl.handle.net/10468/17892

Empowering Explainable Artificial Intelligence Through Case-Based Reasoning: A Comprehensive Exploration

Preeja Pradeep , Marta Caro-Martínez , and Anjana Wijekoon 

(Survey Paper)

I. INTRODUCTION

Abstract—Artificial intelligence (AI) advancements have significantly broadened its application across various sectors, simultaneously elevating concerns regarding the transparency and understandability of AI-driven decisions. Addressing these concerns, this paper embarks on an exploratory journey into Case-Based Reasoning (CBR) and Explainable Artificial Intelligence (XAI), critically examining their convergence and the potential this synergy holds for demystifying the decision-making processes of AI systems. We employ the concept of Explainable CBR (XCBR) system that leverages CBR to acquire case-based explanations or generate explanations using CBR methodologies to enhance AI decision explainability. Though the literature has few surveys on XCBR, recognizing its potential necessitates a detailed exploration of the principles for developing effective XCBR systems. We present a cycle-aligned perspective that examines how explainability functions can be embedded throughout the classical CBR phases: Retrieve, Reuse, Revise, and Retain. Drawing from a comprehensive literature review, we propose a set of six functional goals that reflect key explainability needs. These goals are mapped to six thematic categories, forming the basis of a structured XCBR taxonomy. The discussion extends to the broader challenges and prospects facing the CBR-XAI arena, setting the stage for future research directions. This paper offers design guidance and conceptual grounding for future XCBR research and system development.

Index Terms—Case-based reasoning, explainable artificial intelligence, human-understandable explanations, trustworthy AI, XCBR.

Received 11 March 2024; revised 5 June 2025; accepted 5 September 2025. Date of publication 16 September 2025; date of current version 10 November 2025. This work was supported in part by iSee Project, and in part by CHIST-ERA Project financed by the European Union, for Ireland from the Irish Research Council under Grant CHIST-ERA-2019-iSee (with support the Science Foundation Ireland under Grant 12/RC/2289-P2 at Insight the SFI Research Centre for Data Analytics at UCC, which is co-funded under the European Regional Development Fund), for Spain from the MCIN/AEI and European Union “Next Generation EU/PRTR” under Grant PCI2020-120720-2, with the support of the SHOU-X Project funded by the Comunidad de Madrid through the UCM under Grant PR17/24-31890, and the U.K. from EPSRC under Grant EP/V061755/1. Recommended for acceptance by J. Tang. (Corresponding author: Marta Caro-Martínez.)

Preeja Pradeep is with the Insight Centre for Data Analytics, School of Computer Science and IT, University College Cork, T12 K8AF Cork, Ireland (e-mail: ppradeep@ucc.ie).

Marta Caro-Martínez is with Facultad de Informatica, Universidad Complutense de Madrid, 28040 Madrid, Spain (e-mail: martcaro@ucm.es).

Anjana Wijekoon is with the Department of Computer Science, University College London, WC1E 6BT London, U.K. (e-mail: a.wijekoon@ucl.ac.uk).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TKDE.2025.3609825>, provided by the authors.

Digital Object Identifier 10.1109/TKDE.2025.3609825

IN THE current era of Artificial Intelligence (AI) across various sectors, the pursuit of transparency and interpretability in AI systems has become crucial. Advances in AI, including deep learning and neural networks, have led to significant achievements in healthcare, finance, and autonomous driving. However, the inherent complexity of these models often renders them as ‘black boxes,’ [1] making it challenging to understand their decision-making. This lack of transparency raises biases, ethical integrity, accountability, and fairness [1]. For example, in the justice system, AI assesses recidivism risk scores without clarity on influencing factors. This opacity can result in unfair sentencing, where decisions are based on potentially biased algorithmic outputs rather than the cases’ merits. For users to trust and adopt AI recommendations, they need to understand how decisions are made, enhancing their confidence in the system’s capabilities. For instance, doctors might be reluctant to follow an AI’s cancer treatment recommendation without understanding the underlying rationale. Identifying whether adverse AI decisions are due to system flaws, data bias, or other causes is crucial for legal and ethical accountability. Researchers and practitioners focus on enhancing AI systems’ interpretability and transparency to combat these issues.

One such approach is the integration of Explainable Artificial Intelligence (XAI) and Case-Based Reasoning (CBR). CBR [2] is an intuitive problem-solving paradigm grounded in the cognitive process of reusing past experiences to address new situations through a four-phase cycle: Retrieve, Reuse, Revise, and Retain. New problems are solved by retrieving previously encountered cases similar to the current problem, reusing and adapting their solutions, and then retaining the new experience for future use. CBR, with its analogical reasoning and use of explicit knowledge, offers a naturally interpretable structure that aligns closely with human decision-making processes. However, limited generalization and the need for careful case base maintenance [3] often hinder its utility. In contrast, XAI [4] aims to achieve AI systems transparency using various methods and design principles to provide decision rationales. It provides insights into AI decision-making by promoting responsible AI use; however, it sometimes suffers from oversimplified or misleading explanations that lack personalization. We provide a

TABLE I
 COMPARISON OF OUR XCBR SURVEY WITH EXISTING SURVEYS

Reference	Focus	Description	CBR Cycle Coverage	XAI Coverage	XCBR Taxonomy
Keane et al. [5]	ANN-CBR Twin Systems	Introduce the concept of utilizing CBR as a white-box “twin” for black-box AI systems, including ANN. This twinning approach is a crucial strategy for improving the interpretability of AI models within the XAI framework.	Partial (focus on ANN)	Surface-level (via twin systems)	✗
Weber et al. [9]	XAI in Knowledge-based Systems	Knowledge-based XAI integrates domain knowledge with AI to improve explanations for AI decisions, aligning with CBR principles through supervised data classification. This approach uses AI inputs and outputs as case problems and derives explanation categories as case solutions from domain expertise. Addressing the challenge of accurate classification, it incorporates domain-informed features to refine the selection of explanation categories, thus bridging the data gap between AI operations and human-centric explanations.	No (non-CBR specific)	Moderate (XAI-focused)	✗
Schoenborn et al. [7]	General XCBR Concepts	Highlight the need for explanations in decision-making and present XCBR methods as a solution. Here, XCBR refers to CBR systems that are specifically designed to provide explanations for their outputs. This is distinct from case-based explanations, which utilize CBR as a methodology to elucidate the workings of other systems. It provides a taxonomy of XCBR approaches, making it a valuable resource for XAI and CBR researchers interested in generating and utilizing explanations.	Limited	Minimal	✓(not cycle-based)
Gates and Leake [8]	CBR Explanation Evaluation	Address the need for more assessment in understanding and advancing the explanations of CBR methods for explaining intelligent systems. It proposes strategies for evaluating CBR explanations, surveys XCBR systems, and offers a set of dimensions to categorize the explanation components of CBR systems. It also suggests future research directions and community initiatives to enhance the evaluation and understanding of XCBR systems.	Yes (evaluation only)	Mentioned in evaluation context	✗
Our survey	In-depth XCBR across CBR Cycle	Present a cycle-aligned perspective on XCBR that maps explanatory mechanisms to each phase of the CBR cycle (Retrieve, Reuse, Revise, Retain), thereby clarifying how XCBR systems function, why they support explainability, and what design principles and evaluation strategies are essential for building trustworthy, interpretable AI solutions across high-stakes domains.	Full (Retrieve → Retain)	In-depth (XAI goals, types, evaluation, stakeholders)	✓(cycle-aligned)

concise introduction to the principles and historical development of CBR and XAI in Section III. This paper focuses on Explainable Case-Based Reasoning (XCBR), which synthesizes the strengths of both CBR and XAI. XCBR systems deliver transparent, example-driven explanations grounded in historical data while supporting user-specific insights, adaptability, and continuous learning. XCBR mitigates the limitations of both parent paradigms by combining structured reasoning with explanation clarity, ultimately strengthening user trust and system utility across domains.

Our objective is to examine XCBR from a cycle-aligned perspective, highlighting how existing explanatory mechanisms achieve XAI goals such as transparency, interpretability, and user trust across the CBR cycle. This work aims to clarify how explainability manifests within each CBR phase and provides structured insights for designing, evaluating, and applying XCBR systems in high-stakes domains. An “XCBR system” leverages the inherent interpretability of CBR while embedding explicit explanation mechanisms throughout the reasoning process. It facilitates explainability through garnering case-based explanations [5] or generating explanations following CBR techniques [6]. That means an XCBR system may produce CBR-based explanations by presenting similar past cases to justify its decisions or generate CBR-driven explanations, where the underlying reasoning process is structured and communicated using CBR techniques. While CBR is often considered inherently interpretable due to its reliance on prior cases, XCBR systems go further by explicitly designing to generate and communicate rationale as part of the reasoning process or as an additional explanatory layer. The goal is not merely to retrieve a solution but to justify its selection and explain how it was derived. XCBR utilizes past cases to produce concrete, example-based explanations, similar to how human experts justify their decisions. We will discuss the introduction of XCBR, its taxonomy, and the synergy in more detail in Section IV.

XCBR can be applied in high-stakes domains where transparency, trust, and adaptability are critical. In healthcare, it can support clinicians by retrieving relevant patient cases, explaining treatment rationale, and adapting solutions based on new symptoms or test results. In financial services, it aids fraud detection by comparing transactions with historical patterns and justifying risk scores. Educational tools can use XCBR to recommend personalized learning paths by referencing prior learner profiles and explaining content selection. XCBR systems are also valuable for credit scoring, fraud detection, and risk assessment in banking by adapting decisions from similar past cases. These explanations align with regulatory standards through transparent, example-driven justifications that build user trust.

We reviewed the literature and found a few surveys on CBR for XAI as summarized in Table I. While prior surveys have explored elements of XCBR, our work provides a focused and structured analysis that is unique in several ways. For example, Schoenborn et al. [7] provide a valuable taxonomy of XCBR systems; however, they emphasize classification and do not organize their insights around the complete CBR cycle or pose critical research questions guiding future work. Gates and Leake [8] address explanation evaluation dimensions but do not offer a comprehensive taxonomy. Weber et al. [9] examine XAI in knowledge-based systems without centering on CBR-specific mechanisms. Keane et al. [5] explore CBR as a white-box ‘twin’ to neural networks, but their focus remains on ANN-CBR pairings. In contrast, our work uniquely surveys XCBR from the vantage point of the CBR cycle itself (retrieve, reuse, revise, and retain) and examines how each phase inherently contributes to explainability. Rather than treating XCBR as a modular extension or post-hoc add-on, we emphasize the intrinsic explanatory power embedded in each phase of the CBR pipeline. Our survey addresses pivotal questions and explores the nuanced interplay between CBR and XAI, serving as a foundational guide for researchers in developing XCBR systems. Therefore, our work’s

novelty is elucidating the key processes when using CBR to generate either CBR-based or CBR-driven explanations. We will furnish the specifics of our research method and questions in Section II.

The main contributions of this work are summarized as follows.

- *Cycle-aligned functional perspective on explainability:* We propose a cycle-aligned perspective that highlights how each phase of the CBR cycle: retrieve, reuse, revise, and retain, contributes uniquely to explainability. We identify functional goals from XAI literature through a prospective and integrative analysis that supports these goals within XCBB systems. They are: enhance explainability & semantic interpretation, generate intuitive & relatable explanations, ensure the continual relevance & accuracy of explanations, ensure optimal performance & relevance, refine explanations & improve system performance, and evaluate the performance and explainability.
- *Conceptual foundation for XCBB:* We define XCBB as the integration of CBR's inherently interpretable reasoning with the goals of XAI, highlighting the distinction between CBR-based and CBR-driven explanations. We introduce an XCBB taxonomy that structures the space of explanatory goals and methods, providing a foundation for systematically analyzing and designing explainability within case-based systems.
- *Aligning explainability goals with thematic categories:* We analyze and organize relevant CBR and XAI methods under each functional goal, offering a practical mapping that highlights how these techniques can be used to support the development of explainable capabilities in future XCBB systems. Thematic categories cover aspects such as case representation, semantic interpretation, experience-based reasoning, similarity-based retrieval, adaptation, case base maintenance, and user interaction, which serve as a foundation for guiding the design and evaluation of future XCBB systems.
- *Illustrative use case:* To ground the discussion and demonstrate the practical relevance of XCBB, we adopt an Adverse Drug Reaction (ADR) prediction use case. ADRs are unintended, harmful reactions resulting from the administration of medications in patients receiving polypharmacy, i.e., multiple concurrent medications. Predicting ADR risk is challenging due to complex drug-drug interactions, patient variability, and incomplete clinical knowledge. In this use case, explainability is essential to support clinicians, regulatory authorities, and patients in evaluating model recommendations, highlighting a high-risk decision context.
- *Summary of challenges and future prospects:* We emphasize the challenges in current research and suggest several encouraging avenues for future investigation.

The outline of this article is as follows. Section II provides a concise overview of the research methodology employed in the study, including the resources used and the thought process behind the research development, and ultimately presents the research questions along with the functional explainability goals. Section III discusses the evolutionary trajectory of CBR

and XAI. Afterward, Section IV presents the cycle-aligned perspective on explainability, conceptual foundation of XCBB, and explores how CBR's inherent principles synergize with XAI goals to enhance AI explainability. Section V presents the challenges and prospects in the XCBB arena. Finally, closing remarks are presented in Section VI.

II. RESEARCH METHOD

We undertook a systematic literature review to elucidate CBR's contribution to XAI, aiming to answer six pivotal research questions to pinpoint CBR's specific role within XAI. PRISMA-style flow diagram in Fig. 1 illustrates the methodology for selecting studies in the review. A total of 300 records were initially retrieved from leading academic databases, including Scopus, Google Scholar, IEEE Xplore, ACM Digital Library, and Springer. After removing duplicates and screening titles and abstracts, 250 records remained. Of these, 90 were excluded based on criteria such as lack of peer review, irrelevance to explanation or CBR, or duplication. We assessed 160 full-text articles for eligibility, with 59 excluded due to lack of relevance to XCBB. A final set of 101 articles was selected for in-depth analysis. A noteworthy portion of relevant XCBB research has been showcased at the International Conference on Case-Based Reasoning (ICCBR), including workshops focusing on XCBB. Our systematic review aimed to chart the development trajectory of CBR and XAI, explore their interplay, and highlight the principal challenges and progress made in the field.

To investigate how explainability can be systematically supported in XCBB, we adopt a functional, goal-driven approach that organizes key needs for explainability into seven core dimensions. They are as follows:

- Semantic interpretation [9], [10]: the ability to generate context-aware, conceptually meaningful explanations that align with the user's understanding.
- Intuitive & relatable explanations [11]: the use of relatable, example-based, analogical, or contrastive reasoning that supports user comprehension.
- Continual relevance & accuracy [12]: the capacity for explanations to evolve and remain aligned with changing inputs, knowledge, or user contexts.
- Performance and relevancy [13]: the assurance that explanations remain high-quality, accurate, and appropriate to the task or decision context.
- Explanation refinement [13], [14]: support for interactive or iterative improvement of explanations, informed by user input or feedback.
- Resource efficiency [12]: the ability to scale explanation mechanisms while managing computational or storage costs effectively.
- Evaluation quality [14], [15]: the use of rigorous and multidimensional strategies to assess both the usefulness and trustworthiness of explanations.

Therefore, the six functional goals are: enhance explainability & semantic interpretation, generate intuitive & relatable explanations, ensure the continual relevance & accuracy of explanations, ensure optimal performance & relevance, refine




Fig. 1. Flow diagram for literature selection.




Fig. 2. Mapping of research questions across historical, functional, and future dimensions in XCBR.

explanations & improve system performance, and evaluate the performance and explainability. These goals were identified through an inductive synthesis of foundational and contemporary work on XAI, interpretability, human-centered design, and system performance. Rather than assuming a fixed set of techniques, we frame our research questions around these goals to guide a structured analysis of relevant methods across the literature. This approach allows us to map candidate techniques to each goal, ensuring conceptual clarity and methodological consistency across the cycle-aligned structure of XCBR, which will be presented in Section IV-A. The primary aim of this paper is to provide solutions to the following three main research questions (RQs), which have been formulated based on the literature. Fig. 2 illustrates the mapping of these questions (RQ1–RQ3) across historical, functional, and future dimensions in XCBR.

- RQ1:** What is the historical progression of CBR and XAI as distinct fields?
 The motivation is to understand how foundational ideas in both CBR and XAI originated and evolved, providing context to current practices and methodologies, which will be explained in Section III.
- RQ2:** How do CBR's inherent principles synergize with XAI's goals to enhance AI explainability, and what key

features or functionalities are critical in this integration of strengthening explainability across varied user expertise? This investigation examines how CBR's inherent principles align with XAI goals to enhance explainability and identifies critical features and functionalities across the Retrieve, Reuse, Revise, and Retain phases. This supports clear, relevant, and user-adaptable explanations with varying levels of expertise, which will be discussed in Section IV. In particular, we are interested in reviewing the following.

- RQ2.1:** How can XCBR systems enhance explainability and semantic interpretation?
- RQ2.2:** What techniques can be used to generate intuitive and relatable explanations in XCBR systems?
- RQ2.3:** What are the effective strategies for adaptation and learning in XCBR systems to ensure the continual relevance and accuracy of explanations?
- RQ2.4:** How can efficient case storage and management support optimal performance and long-term relevance in XCBR systems?
- RQ2.5:** What techniques can incorporate user interaction and feedback to refine explanations and improve system performance in XCBR systems?




Fig. 3. Core concepts of CBR including representation, reasoning cycle, and knowledge containers.

- *RQ2.6*: What strategies and metrics can be used to evaluate the performance and explainability of XCBR systems?
 - *RQ3*: What are the current challenges facing the CBR-XAI arena, and what potential research avenues can address these to drive future advancements?
- This identifies the prevailing challenges within the CBR-XAI domain and explores potential research directions that could address these challenges, fostering advancements in the field, which will be explored in Section V.

III. HISTORICAL PROGRESSION OF CBR AND XAI

In this section, we address the RQ1 in Section II by providing the evolutionary trajectory of CBR and XAI domains. CBR and XAI have rich histories and evolved, responding to technological advancements and shifting needs in AI. This section seeks to understand how each field has evolved from its inception to its current state, including its foundational theories and technological advancements.

A. Progression and Core Concepts of CBR

In the 1980 s and 1990 s, CBR emerged as a notable AI methodology, emphasizing problem-solving by reusing past cases or instances [16]. Its evolution from a cognitive model to a robust AI methodology underscores its adaptability and enduring significance. CBR continues to adapt and incorporate technological innovations, expanding its application across diverse fields. This section delves into the core components of CBR, demonstrating how it utilizes historical experiences to address new problems. We focus on three key areas (as illustrated in Fig. 3): the advancements in case representation within CBR systems, the analysis of the CBR cycle, and the evolution of knowledge containers.

1) *Case Representation*: In CBR, knowledge is encapsulated in cases comprising a problem description, the solution applied, and the solution's outcome [2]. The problem description details the context, while the solution addresses the problem. These cases represent real-world scenarios and solutions, stored in a case base that can vary in format, including databases, text files, or XML files [17]. Addressing the complexity of design cases necessitates enhancing cases with generalized design knowledge and formalizing typically informal knowledge [18]. Case representation involves the structuring and storage of cases. Various

formalisms for case representation include feature vector, structured, and textual formats [17], [19]. In *feature-vector representation*, cases are defined by attribute-value pairs, allowing for straightforward comparisons but needing more domain knowledge and semantic depth. Meanwhile, *structured representations* manage cases in a frame-based structure or relational representations by considering cases as clusters of related objects similar to episodic memory. *Textual case representation* uses text to describe cases and supports problem-solving by comparing textual descriptions of past cases. There exist different methods for case representation, including ontologies [20], [21], XML for marking up cases [22], and knowledge graphs [18] that offer a range of approaches to structuring and managing cases in CBR systems.

2) *Four-Phase CBR Cycle*: The CBR cycle is a constantly evolving process enriched with every new problem and solution, which relies on historical data and previous experiences to facilitate iterative and hands-on problem solving. It is a practical and intuitive approach that finds widespread application in various domains, including medical diagnosis, legal reasoning, and customer support [23]. The CBR cycle typically consists of four main phases as follows [16], [24].

- *Retrieval*: The CBR system calculates similarity scores between a new and unseen instance and all stored cases in the case base. The system uses predefined similarity metrics to compare features, attributes, or domain-specific criteria. The system then retrieves the most similar past cases based on these similarity scores. These cases have historically dealt with problems most akin to the new instance, so they serve as the primary candidates for explanations. For instance, a CBR system for medical diagnosis retrieves cases of patients with similar medical histories and symptoms. Ensuring quick retrieval is vital to defining the problem accurately and applying a relevant similarity metric to discover the most pertinent past cases.
- *Reuse or Adaptation*: The CBR system might adapt the solution from the past case to better align with the current scenario, where the retrieved case is not a perfect match to the new problem. CBR reuses the solutions or parts of the solutions from similar cases to solve the new problem after retrieving the relevant cases. This outcome is based on past experiences and is adapted to meet current needs by providing a clear explanation. For instance, a successful treatment plan from the past can be modified for a new

patient by considering their unique characteristics and any updates in medical practice.

- *Revision*: The CBR system evaluates the outcome after solving the new problem. If necessary, the system revises the case by refining the proposed solution from the reuse phase. The CBR system’s performance and efficacy improve through this continuous learning process over time. For example, it keeps track of the patient’s reaction to the customized treatment plan and makes necessary adjustments.
- *Retain*: This involves updating the case base. At this point, the system stores the new case that was just solved, including the problem description, the adapted solution, and any additional information that might be useful for future problem-solving. For instance, documenting the new patient case, including the adapted treatment and outcomes, can assist with future medical diagnoses.

3) *CBR Knowledge Containers*: Implementing specific functions aimed at uniformly describing problems, ensuring equitable comparisons, and facilitating adaptation for generating feasible solutions is crucial at various stages of a CBR system. These functions are encapsulated as ‘knowledge containers’ [25], which model the knowledge within a CBR system. A CBR system can be conceptualized by interacting with four primary knowledge containers: vocabulary knowledge, similarity knowledge, adaptation knowledge, and case base knowledge [25], [26], [27]. *Vocabulary knowledge* encompasses the terminologies and descriptions used for cases, whereas *Similarity knowledge* involves the criteria for assessing case similarities. *Adaptation knowledge* is used during the CBR cycle’s revision phase and focuses on adjusting solutions from the most similar past cases to fit new problems. Meanwhile, *case base knowledge* highlights the case base’s scope and identifies areas for expansion to enhance future case acquisition by reflecting the problems in the stored cases.

It is uncommon for all knowledge containers to be fully developed due to the costs and complexities involved. Typically, some containers are less developed while others are more advanced, maintaining the system’s overall performance. Distinct algorithms usually drive each container’s functionality; for example, the similarity container might use semantic knowledge from an ontology in a textual CBR system, while the vocabulary container relies on complex text features [27].

B. Evolution and Key Fundamentals of XAI

Early AI systems were inherently interpretable due to hand-crafted rules for decision-making on raw data. The rise of black-box models, on the other hand, made the need for explainability much greater, which led to the growth of the XAI field. Recent research has delved into XAI’s various facets, including ‘what,’ ‘why,’ ‘what for,’ and ‘how’ [12], [23], [28]. Understanding the target audience, the reasons behind needing an explanation, and the context in which it is provided are all necessary for effective explainability. Considering these factors, explanations can be crafted more meaningfully and user-centric, presented at the right time, and in the most suitable format.

We need to clarify the term “explainability” as the literature [23], [28] often uses related terms, including “interpretability,” “transparency,” “comprehensibility,” and “understandability.” Although these terms share a similar meaning, their undertones differ slightly. Explainability refers to making AI system operations understandable to users by providing interpretable explanations that bridge system decisions with user comprehension, such as creditworthiness reports explaining loan approvals based on transparent rules. Interpretability emphasizes how a system conveys its actions even to non-experts, for instance, AI dietary recommendations clarifying suggestions based on health data. Transparency highlights a model’s inherent self-explanation ability without extra tools, as seen in real estate price predictions where valuation logic is evident. Comprehensibility presents model knowledge in human-like reasoning formats, exemplified by medical diagnosis explanations following familiar rule-based logic. Understandability centers on enabling users to grasp AI’s purpose and conclusions without requiring deep technical knowledge, such as AI-driven financial planning that adapts advice to individual goals while remaining accessible. Generating explanations can demand substantial resources and may not always be necessary, though it becomes critical in safety-sensitive contexts where errors could lead to severe consequences. In the following sections, we will examine the role of explainability in AI, determine who needs explanations, discuss the goals of XAI pertinent to its target audience, and discuss the diverse methods, levels, and strategies for crafting explanations in XAI systems. Our discussion also aims to showcase the efforts within the field to render AI decisions comprehensible and responsible.

1) *Why XAI?*: Explainability in AI serves multiple roles based on specific requirements [28], going beyond just improving predictive performance. For instance, explainability helps evaluate loan applications in banking and finance and sheds light on the factors affecting the AI model’s decisions. Explainability reveals the factors influencing the model’s decision, enabling analysts to confirm its legality and fairness and make necessary adjustments. This aspect is vital for adhering to legal standards, such as the European Union’s General Data Protection Regulation (GDPR)’s ‘right to explanation’ [29]. Explainability helps clinicians justify their treatment recommendations to patients, bolstering trust by facilitating informed healthcare decisions. It also improves the learning and knowledge discovery of AI models. This enables the comparison of AI-derived strategies with established knowledge for educational and scientific advancements. Explainability in AI models helps the pharmaceutical industry find new drug candidates by clarifying how the AI predicts therapeutic effects or drug interactions.

XAI is critical in fostering trust, accountability, and adherence to regulatory standards in AI systems. It enhances trust by making decision-making transparent and vital in healthcare, finance, legal sectors, and autonomous systems. Making a system’s decision-making process more transparent increases user confidence, promotes adoption, and reduces skepticism about AI. Although some studies emphasize the importance of explainability in fostering trust, it is essential to recognize that trust is not the sole aim of explainable models; various goals can

coexist [23]. XAI also underpins accountability [30], allowing users to evaluate an AI system's fairness, biases, and ethical considerations. It aids stakeholders in ensuring AI systems meet legal, ethical, and regulatory standards by shedding light on the decision-making criteria and identifying and correcting biases or discriminatory practices. Additionally, XAI is essential for complying with regulations such as the GDPR, which emphasizes the need for individuals to understand the logic, importance, and repercussions of automated decisions affecting them [4]. XAI enables organizations to offer clear explanations and rationales for AI-based decisions, helping them comply with regulatory demands, sustain public trust, and prevent legal issues.

2) *Target Audience/Stakeholders in XAI*: XAI is an AI field that can be applied in any domain and has to satisfy their users in those domains. These stakeholders represent a broad spectrum of expertise, including data scientists, domain experts, regulators, end-users, developers, and more, each contributing unique perspectives and requirements to the XAI field [30]. Understanding the motivations and needs of these critical groups is crucial for crafting effective XAI solutions that meet ethical, legal, and domain-specific standards. *AI Experts (Data Scientists and Machine Learning (ML) Practitioners, AI System Developers)* drive AI model development by understanding complex behaviors, identifying improvement areas, and refining models, with goals of gaining insights, ensuring fairness, and resolving prediction issues using XAI tools. *Domain Experts* ensures reliable advice and regulatory compliance, healthcare, finance, and engineering experts align AI models with industry standards and ethics. *Regulators and Legal Authorities* monitor AI implementations to comply with legal and ethical guidelines, using XAI to evaluate system transparency, fairness, and adherence to regulations such as GDPR and HIPAA. *End-Users or Consumers* demands transparent AI explanations to make well-informed decisions. Finally, *Ethics Committees and Review Boards* evaluate AI applications for ethical integrity, protocol compliance, and patient safety to ensure decisions prioritize well-being.

3) *XAI Goals and Stakeholder Relevance*: The pursuit of explainability in AI is not just a technical challenge but a multifaceted endeavor to address the specific needs of various stakeholders. XAI aims to provide transparency and understandability to AI systems by fostering trust and facilitating responsible usage. Research shows that the main goals of XAI [21], [23] are attuned to the informational needs of distinct stakeholders. From now on, we will use the term 'stakeholders' to broadly encompass all parties involved in system design, development, oversight, and usage.

- 1) *Informativeness*: Provides detailed insights into the model's functioning to support informed decision-making across all stakeholder groups. For example, an AI predicting diabetes risk explains the impact of blood glucose levels, facilitating targeted health interventions.
- 2) *Transferability*: ML models must adapt across diverse scenarios while acknowledging limitations. Explainability aids stakeholders, namely domain experts and data scientists, in applying insights across different domains. For instance, a company uses a recommendation model for cultural and geographical variations in a new market.

- 3) *Accessibility*: Simplifies complex algorithms for non-experts, enhancing user-friendliness. Key stakeholders, mainly end-users or consumers, benefit directly, as seen in healthcare applications where doctors receive understandable AI diagnoses to foster trust and improve patient care.
- 4) *Fairness*: Identifying biases in data ensures stakeholders, such as regulators, legal authorities, and end-users or consumers, get equitable outcomes by allowing corrections. XAI reveals biases in AI recruitment systems, namely age or gender disparities, enabling organizations to correct unfair practices for a more equitable hiring process.
- 5) *Confidence*: Confidence in AI is critical for reliability and entails robustness and stability. Explainable models offer confidence levels vital for stakeholders, including developers, end-users, and regulators. For example, an AI in autonomous vehicles that provides a confidence score for its decisions, namely braking, enhances transparency and safety.
- 6) *Usability*: Engages stakeholders through feedback and customization, improving trust and interaction. An investment platform, for example, uses AI to tailor advice based on user-defined risk preferences and investment goals.
- 7) *Causality*: Identifying potential causal relationships differentiates correlation from causation, validating AI outcomes for stakeholders, including domain experts and regulatory authorities. An autonomous vehicle company may link rapid acceleration to higher accident rates through controlled experiments, thereby enhancing vehicle safety.

C. Type, Scope, and Techniques for Explanation

Classifications of XAI techniques focus on their explanation mechanisms, types, and scopes, aiming to clarify AI decision-making for various data types [12], [31]. To ensure AI decisions are understandable, it is crucial to tailor explanations to the data type — tabular, image, or text [31]. *Tabular data* is structured and can be used to generate logical explanations that are easy for stakeholders to understand. We can apply decision trees to tabular data to create rule-based explanations that are transparent and straightforward. For *image data*, explanations often come in visualizations highlighting an image's most influential areas in the model's decision-making process, allowing stakeholders to understand why a model made a particular classification or recognition decision. In the case of *text data*, explanations can be provided by identifying and presenting the key phrases or words that influenced the model's predictions.

The literature identifies nine types of explanations: case-based, contrastive, counterfactual, trace-based, feature importance, anchor, saliency map, factual, and semi-factual, each suited for different AI scenarios. *Case-based explanation* [5] uses historical solutions to solve current problems, similar to how a doctor might recommend treatments based on past successes with similar symptoms. Another type is *contrastive explanation* [32], which highlights the differences leading to a specific outcome, as seen in a loan approval system that contrasts credit histories to explain different decisions. *Counterfactual*

explanation [32], [33] explores alternate scenarios by altering inputs, such as a traffic navigation app suggesting an earlier departure to avoid traffic. Another method, *trace-based explanation* [34], documents the sequence of steps the system takes to reach a conclusion, such as a diagnostic AI explaining the tests conducted for a diagnosis.

Feature importance [35] explains model decisions by quantifying the contribution of each input feature to the output, aiding stakeholders in understanding influential variables. For example, a credit scoring model shows that income contributes 40% to approval, while debt ratio contributes 25%. *Anchors* [36] are if-then rule-based explanations that identify crucial feature conditions supporting a prediction. For instance, an email classifier explains that if the email sender is unknown and the subject contains ‘Congratulations,’ the message will be classified as spam. *Saliency maps* [37] indicate which parts of input data (often image pixels) influence the output, such as highlighting areas of a lung scan that affect a pneumonia diagnosis. *Factual explanation* [38] clarify why a specific prediction was made by offering supporting evidence, such as an approved loan due to the applicant’s income and credit score meeting standards. Finally, *semi-factuals* [39] explain what would have happened if specific inputs had been slightly different but still resulted in the same decision. For example, the applicant’s outstanding credit history would have allowed the loan to be approved even if their income had been somewhat lower.

We categorize explainability techniques into portability and concurrentness based on the iSee framework [20]. Portability describes how dependent an explanation technique is on the underlying model architecture, distinguishing between model-specific or model-agnostic. *Model-agnostic* method explains black-box models by approximating behavior from inputs and outputs (e.g., SHapley Additive exPlanations (SHAP), Local Interpretable Model-Agnostic Explanations (LIME). Meanwhile, *Model-specific* focuses on details of a particular model (e.g., Deep Learning Important FeaTures (DeepLift), or Integrated Gradients). Another approach is the influence function [40], [41], which estimates the effect of removing or upweighting a single training instance on the prediction for a given test instance. Influence functions, a concept from statistics, have been adapted for use in XAI to analyze the impact of individual training examples on model predictions [40], [41]. They estimate how the model’s output for a specific test instance would change if a training instance were unweighted or removed. Although this method does not provide the same type of explanation as CBR, which returns the most similar past cases to the current one, the two approaches complement each other. Both focus on linking a prediction to specific examples and thus could be used complementary to generate more enriched and multi-faceted explanations [40], [41].

Another dimension is concurrentness, which describes when explanations are produced relative to model training, categorized into ante-hoc or post-hoc techniques. Ante-hoc methods integrate explainability directly into the model during its construction, producing inherently interpretable models. In contrast, post-hoc techniques generate explanations after training the model, often applied externally to black-box models. Model-agnostic explainers operate independently of model architecture

and are usually post-hoc, while model-specific explainers are tailored for specific model types (e.g., CNNs). When model-specific explainers use internal model information, they are ante-hoc; otherwise, they are post-hoc. Ante-hoc approaches can sometimes be model-agnostic, particularly in interpretable models like rule lists. Overall, model-agnostic and model-specific refer to the explainer’s dependence on the model, while ante-hoc and post-hoc indicate when explanations are integrated.

Post-hoc explainability techniques elucidate non-intuitive models via a range of methods, including textual and visual explanations, as well as explanations by feature relevance, example, and simplification, each addressing different facets of explainability, making complex AI transparent [23]. *Text explanations* demystify model processes with understandable language, while *visual explanations* use graphics to simplify and clarify model behavior. *Local explanations* dissect decisions into smaller parts for detailed understanding, whereas *Feature relevance* clarifies how specific data aspects influence outcomes. Meanwhile, *explanations-by-examples* uses data instances to show decision impact akin to human learning. Additionally, *explanations-by-simplification* improves stakeholder comprehension of AI mechanisms by producing a simpler model replicating the original’s key features. Distinguishing between an explanation’s format and type is essential. Reasoning-based forms, namely contrastive, counterfactual, and trace-based explanations, should be grouped alongside some explanation types, such as feature relevance, example-based, and simplification-based explanations. These represent different underlying mechanisms or strategies for explanation. However, visual and textual explanations refer to the presentation format of these underlying explanation types and should be discussed separately regarding the user interface and communication.

Explanation scope defines the extent of interpretation: it can be *global*, covering the entire model, or *local*, focusing on the rationale behind specific predictions [30], [31]. Local explanations aim to clarify the reasoning behind a single prediction, whereas global explanations provide insight into the model’s overall behavior. Notably, local explanations are not always post-hoc. For instance, SHAP is a post-hoc explainer that can be used to generate both local and global explanations, based on its application.

In the XAI literature, various inherently explainable models have been developed to ensure interpretability by design. These include linear models, generalized additive models, decision trees, rule-based systems, scorecards, and neural additive models [23]. Each model type carries specific assumptions and explanation capabilities. Numerous libraries for explanation methods are also available, such as IBM Research’s AI Explainability 360 [42], and Facebook’s Captum [43]. The iSee: Intelligent Sharing of Explanation Experience platform¹ also includes an Explanation Library, which amalgamates over 70 explainers from assorted XAI libraries, accessible via its GitHub repository². It includes unique explainers developed by the iSee team, such as DiSCERN [44], and PertCF [45].

¹ <https://isee4xai.com/>

² <https://github.com/isee4xai/iSeeExplainerLibrary>




Fig. 4. Conceptual integration of CBR and XAI into XCBR.

IV. CBR AND XAI SYNERGY: TOWARD XCBR

In this section, we tackle RQ2 in Section II by examining the synergy between CBR and XAI principles to enhance explainability. It explores how CBR's approach of using past cases to solve new problems complements XAI's goals for transparency, examines the impact of combining CBR with XAI to create more understandable AI systems, and identifies essential features enabling their integration. As Section III-B suggests, tailoring explanations to varied stakeholder needs and contexts is crucial yet challenging. A unified platform could simplify this by enabling the reuse of optimal explanation strategies across different scenarios.

A. Cycle-Aligned Perspective on XCBR

Although CBR contributes interpretability to AI systems by design and has been applied across several explanation types, these applications are often informal, fragmented, or manually driven. XCBR is a systematic, user-centered framework that formalizes and extends the explanatory potential of CBR within the broader XAI landscape. Fig. 4 shows how XCBR arises from the intersection of CBR and XAI. The strengths of CBR and XAI are combined in XCBR to provide user-centered explanations based on real-world examples. We structure XCBR according to the CBR cycle and tailor it to the stakeholder's mental model, where each phase contributes distinct explanatory opportunities, such as retrieval rationales, reuse justifications, revision insights, and retention summaries. XCBR unites the interpretability and analogy-driven decision-making of CBR with the transparency, accountability, and human-centered goals of XAI. The result is an AI paradigm that solves problems by referencing past cases and explains its reasoning in an understandable, relevant, and actionable way. Effective XCBR systems design requires more than leveraging analogy; it also needs careful consideration of the explanatory requirements and design at each stage of the reasoning process. While CBR and XAI each offer inherent forms of interpretability, XCBR formalizes the design of systems that




Fig. 5. Taxonomy of XCBR by organizing explainability approaches for CBR into four major categories: Explanation Type, Explanation Generation, Explanation Modality, and Integration.

reason from cases and explain that reasoning in a transparent, stakeholder-centered manner. In this section, we will discuss the functional components and taxonomy of XCBR along with an illustrative use case.

1) *Functional Components of XCBR*: XCBR systems enhance transparency, stakeholder trust, and decision accountability by incorporating explicit explanation mechanisms into each phase of the CBR cycle (retrieval, reuse, revision, and retention). XCBR highlights the relevance of features and justifies the similarity measures used to select past cases during case retrieval. The system explains how the retrieved solution is modified to fit the new problem context in the reuse/adaptation phase. Furthermore, the revision step traces any corrections made to the adapted solution, offering justifications for adjustments based on domain feedback. Finally, in the retention phase, XCBR systems provide reasons for storing the new case guided by novelty, utility, or diversity criteria.

2) *Taxonomy of XCBR*: We present a taxonomy that classifies systems along four key dimensions, as shown in Fig. 5 to better understand the diversity of approaches within XCBR. While the taxonomy presented by Schoenborn et al. [7] offers a foundational structure to classify contributions in XCBR, it remains broad and centered around conceptual categories. In contrast, our taxonomy is grounded in the functional design of XCBR systems, emphasizing how explanations are generated, presented, and integrated within the CBR cycle. This approach allows for a thorough understanding of how XCBR systems reason, explain, communicate, and interact with other AI frameworks. Researchers and developers can better assess suitability for specific applications, identify gaps, and design more effective explainable reasoning frameworks by categorizing systems along these four dimensions. Moreover, the proposed taxonomy reflects emerging trends in XAI and stakeholder-centered explanation, offering a practical framework for system development, evaluation, and cross-domain comparison. The use of XCBR spans multiple domains, such as healthcare (e.g., diagnostic support), legal reasoning (e.g., precedent analysis), recommender systems (e.g., personalized suggestions), finance (e.g., risk assessment), and education (e.g., adaptive tutoring), each influencing the type of explanation required.

- *Explanation Type*: XCBR systems typically offer two main categories of explanations: CBR-based and CBR-driven, as stated in Section I. *CBR-based explanations* justify decisions by directly referencing similar past cases from the case base. Drawing analogies with familiar or previously successful situations helps stakeholders understand decisions. These explanations include narrative comparisons,

side-by-side visualizations, or annotated case matchings. *CBR-driven explanations*, on the other hand, go beyond mere retrieval by using CBR techniques to shape or construct the explanation itself actively. These systems utilize analogical reasoning, case adaptation traces, or structured argumentation models to generate coherent, context-sensitive justifications that reflect expert reasoning. iSee platform is a CBR-driven XAI recommender focused on improving AI system explainability through advanced abstraction in explanations [6]. Through an Explanation Experiences Editor (iSeeE3), AI designers can capture and disseminate complex explanation experiences among peers with similar needs.

- *Explanation Generation*: The way the explanations are generated, whether static or dynamic, is a crucial aspect of the XCBR taxonomy. *Static explanations* are pre-defined during system design or case base construction. These explanations do not adapt to stakeholder queries or context; they are typically associated with specific cases or solutions. Although practical and straightforward to implement, they frequently lack personalization and might not meet the stakeholder's urgent informational needs. *Interactive explanations* are generated at runtime in response to stakeholder queries, contextual variables, or the system's confidence in its solution. These explanations can include additional reasoning traces, emphasize feature importance, or adjust content based on stakeholder interactions. Even though they might require additional computational resources, interactive explanations tend to be more flexible, interactive, and cognitively aligned with stakeholder expectations. Relatively little research has been done on interactive XCBR systems. In contrast, most current literature on XCBR concentrates on static explanations, in which arguments and examples are given without being adjusted for the stakeholder's input or context. Stakeholder interaction in interactive XCBR generates and refines explanations by incorporating real-time feedback, preferences, or clarification requests.
- *Explanation Modality*: XCBR systems differ significantly in how explanations are presented to stakeholders. The modality of explanation affects interpretability, cognitive load, and stakeholder engagement. Modality can be customized according to stakeholder preferences, domain complexity, and the purpose of the explanation, whether it is to justify a decision, support learning, or facilitate trust in the system's recommendations. The explanation modalities in XCBR are:
 - *Textual explanations* are the most common modality, often delivered as natural language justifications, rule-based statements, annotated case narratives, or even as tables. They are effective in healthcare and legal reasoning domains, where domain-specific language is critical. In XCBR, the textual modality presents information about previous cases in natural language as explanation cases. The goal is to offer the stakeholder a better understanding of the knowledge provided within the explanation cases.

- *Visual explanations* use graphical representations to convey reasoning, such as similarity maps, adaptation graphs, feature heatmaps, plots, trees, annotated images, or side-by-side comparisons. Visual modality is beneficial in image-based systems or when summarizing complex case relationships. In XCBR, including visual elements is essential for improving the comprehension of the provided explanation cases, particularly when the AI model utilizes images as part of its data.

- *Integration with XAI Techniques*: Modern XCBR systems vary in how much they integrate with broader XAI frameworks. We classify them as pure, twin, and hybrid systems. *Pure XCBR systems* refer to CBR-based systems that generate explanations alongside predictions in an ante-hoc manner. In these systems, predictions are produced using CBR methods, and explanations are derived through the exact CBR mechanisms, typically by retrieving and adapting relevant past cases to justify the current decision. *Twin XCBR systems* generate explanations using CBR mechanisms such as similarity metrics, retrieval justifications, and adaptation traces. They usually do not use external explainability tools but rely on rule-based or symbolic representations. The key distinction from pure XCBR systems is that, in twin XCBR, the predictions being explained are produced by a separate AI model, often a neural network. The CBR component functions solely as an explanation module, interpreting and explaining the behavior of the underlying model. 'Twin-ing,' a black-box AI with a CBR twin proposed by Keane and Smyth [5], creates an interpretable proxy using CBR retrieval to explain decisions with similar past examples. *Hybrid XCBR systems* employ model-agnostic XAI methods to improve their reasoning or explanations, such as LIME, SHAP, Anchors, or Counterfactuals. These integrations provide additional transparency, primarily when the CBR system interacts with black-box models (e.g., neural networks) by facilitating multi-perspective explanations, such as combining similarity reasoning with feature attribution. For instance, Tofighi et al. [46] combined CBR with MLP classifiers and employed SHAP-based feature attributions to explain lung cancer diagnosis decisions. The level of integration affects the system's interpretability, complexity, and explainability across use cases. Our 'integration' dimension focuses on how CBR systems interface with external, model-agnostic XAI techniques. While substantial work is on pure XCBR and twin XCBR systems, research on hybrid XCBR remains relatively limited. These approaches are vital as they offer the potential to combine predictive strength with diverse explanatory strategies. In user-facing applications, different stakeholders might need various perspectives or levels of detail in the explanations provided. To meet the varied interpretability needs of real-world decision-making contexts, hybrid XCBR research must be expanded.

3) *ADR Prediction Use Case*: We present the ADR prediction use case discussed in Section I to illustrate the functional




Fig. 6. Hierarchical view of XCBR based on core CBR cycles and their specific contributions to explainability.

components of XCBR. Accurate ADR risk prediction models are required for enhancing patient safety, optimizing treatments, and minimizing preventable medical complications. For clinicians to understand the reasoning behind predictions, such as risk factors or drug combinations, explainability is essential. Transparent explanations build trust, ensure regulatory compliance, and promote ethical accountability, making this area ideal for evaluating explainability and CBR methods.

To illustrate the taxonomy, we consider a clinical decision support system developed for the ADR prediction use case. The system provides CBR-based explanations by retrieving similar past ADR cases based on drug combinations and lab values. CBR-driven explanations show how the retrieved recommendations are adapted to patient-specific contexts such as impaired renal function or age. The system dynamically generates real-time explanations with justifications tailored to the clinical scenario and stakeholder role (e.g., physician or pharmacist). It offers textual summaries of clinical reasoning (e.g., “ADR risk increased due to interaction between drug X and elevated creatinine”), visual timelines of drug administration with flagged interactions, and interactive tools for exploring medication changes. Finally, it demonstrates hybrid integration with XAI techniques, e.g., SHAP, to identify which features most influenced the prediction, thereby reinforcing transparency and stakeholder trust. In model-agnostic setups, SHAP can explain any black-box model by approximating feature attributions. In addition, we can use influence functions to identify which specific training cases have most significantly shaped the model’s prediction, providing a causally grounded, data-centric explanation that complements the similarity-based reasoning of CBR. This enables stakeholders to understand the features that contributed to the prediction and the past cases in the dataset that were most responsible for forming the decision boundary relevant to the current case.

B. Aligning Explainability Goals With Thematic Categories

We define a set of research questions aligned with key functional goals for explainability in XCBR to guide our analysis. As discussed in Section II, the six functional goals identified are to enhance explainability & semantic interpretation, generate intuitive & relatable explanations, ensure the continual relevance

& accuracy of explanations, ensure optimal performance & relevance, refine explanations & improve system performance, and evaluate the performance and explainability. Each question targets a specific explanatory objective and supports identifying relevant techniques from CBR and XAI literature. Fig. 6 illustrates a CBR cycle-aligned hierarchy of XCBR developed through a comprehensive literature synthesis. This hierarchy is organized around the four canonical CBR phases, each annotated with core functionalities supporting explainability. In our previous paper [47], we discussed essential CBR principles and their relevance to XAI. The literature provides empirical and theoretical justification for including each functionality mentioned in Fig. 6. El-Sappagh & Elmoghy [17] underscore the value of structured *case representation*, which supports traceable decision-making by enabling explicit mappings between past instances and present outputs. Bergmann, Pews, & Wilke [48] argue for the importance of *domain knowledge* integration, which enriches the semantic grounding of explanations to offer deeper, context-aware insights. Cañas, Leake, & Maguitman [49] state that *experience-based reasoning* fosters intuitive understanding by drawing from analogous cases. Similarly, De Mantaras et al. [24] emphasize *similarity-based retrieval* as a foundation for objective, transparent case selection. To refine retrieved solutions and supporting system learning, Wilke & Bergmann [26] emphasize *Adaptation and learning*, which is central to the reuse phase. *Case base maintenance (CBM)* discussed by Chebel-Morello, Haouchine, and Zerhouni [3] ensures that explanations evolve by removing obsolete cases and integrating new and validated experiences. Over time, this will reinforce trust and accuracy. Finally, Sokol & Flach [50] highlight CBR’s *interactive and iterative* nature, where stakeholder dialogue enables clarification, validation, and personalization of explanations.

Furthermore, the hierarchical structure was informed by considering how different XCBR systems embed explainability in the reasoning cycle and translating these findings into critical research questions revealed in this study. One cross-cutting factor, *Strategies & Metrics*, address top-down issues involving multiple phases. These dimensions encompass continuous learning, stakeholder-centered evaluation, and system efficiency. This framework offers a methodical and comprehensive viewpoint because the taxonomy is rooted in the CBR cycle. It guides the

design, analysis, and implementation of transparent, adaptive, and stakeholder-aligned XCBR systems.

Our literature review identifies key categories vital for CBR and XAI synergy, which address RQ2.1 to RQ2.6 in Section II. In subsequent sections, we discuss how these elements enhance XAI systems' transparency, understandability, and interpretability. In each section, we will also discuss the ADR prediction use case to provide more insight on how to incorporate the concepts.

1) *RQ2.1. Enhancing Explainability and Semantic Interpretation Through Case Structuring and Domain Knowledge:* In this section, we examine techniques from the CBR literature focused on enhancing case representation with domain knowledge to improve semantic transparency, stakeholder understanding, and contextual relevance. Ontologies and domain knowledge enrich case representation, enhance semantic retrieval, and guide explanation generation in XCBR systems. CBR enhances XAI by providing relatable examples, making AI decisions tangible and understandable through real-world instances. This approach contrasts with abstract, technical explanations and grounding decision-making in similar past cases. The effectiveness of CBR hinges on precise case representation, as discussed in Section III-A1, within the system's knowledge base, which is crucial for matching new problems with accurate solutions. A clear, comprehensive case representation is essential for offering understandable explanations to stakeholders in an XCBR system [51].

Gaining a deep insight into case representation in XCBR is crucial, which can be accomplished by examining the functions, challenges, and considerations highlighted in Section III-A1. Incorporating domain knowledge [9] refines explanations, making AI decisions more transparent and relevant, especially in critical sectors such as finance and healthcare. This method utilizes expert insights and particular data, enhancing trust and supporting better decision-making. Domain-specific explanations help stakeholders understand the rationale behind AI recommendations, while post-hoc verification ensures their scientific validity, fostering collaboration and improving real-world outcomes. Granularity in representation is vital for providing clear, tailored explanations in XCBR systems, meeting the varied needs of stakeholders from experts to laypersons. This approach enhances stakeholder trust, improves understandability, and facilitates decision-making. Moreover, addressing the complexities of Big Data [52], XCBR systems must adeptly handle various data types, ensuring process transparency and comprehensible data representation for stakeholders. Precise feature extraction and adaptable representation methods are crucial for maintaining transparency and enhancing system trust and efficiency. The structural consistency and semantic depth in case representation play vital roles in XCBR systems. They ensure the system's usability and comprehensibility, facilitating stakeholder engagement by providing context-rich, meaningful explanations that align with stakeholder needs and preferences, thus bolstering trust and decision-making in XAI applications.

We will investigate how ontology enhances CBR systems, improving case indexing, retrieval, representation, and more [17]. By representing case bases as ontologies, it eases knowledge

acquisition and case discovery. As frameworks detailing domain knowledge, ontologies define concepts and relationships [53] vital for AI, ensuring cases are well-described and domain-specific [54]. This clarity and semantic richness allow for precise case comparisons using semantic similarity metrics [55]. Machine-readable ontologies support the creation of diverse, user-specific explanations [56], as demonstrated by several studies. Tiddi et al. [57] devised an Ontology Design Pattern (ODP) tailored for explanation representation, illustrating its capability to encapsulate explanatory concepts across diverse disciplines. A formalized domain-specific model, Food Explanation Ontology (FEO) [58] integrates concepts from the explanation domain to formulate answers to stakeholder inquiries regarding food suggestions offered by AI systems, such as personalized knowledge base question-answering platforms. Another proposal [21] designed a conceptual model and an ontology, ExRecOnto, to integrate explanations into recommender systems, focusing on stakeholder-centric, practical explanations.

iSeeOnto [20], created within the iSee framework, identified crucial concepts for describing explanation experiences in a CBR-driven XAI. iSeeOnto case description consists of attributes such as AI Task, AI Method, Dataset Type, Portability, Scope, Target, Presentation, Concurrentness, Intent, Technical Facilities, AI Knowledge, Domain Knowledge, and User Questions [6]. More details of iSeeOnto can be found here³. Chari et al. [56] created an explanation ontology to model essential elements for crafting user-focused explanations, covering various explanatory questions 'How, Why, Why-not, What-if, and How-to.' Domain-specific knowledge from multiple fields has been applied to enhance black-box model explainability without compromising performance [10]. Integrating this knowledge with ontological structures into XCBR enhances solution clarity and accuracy, advancing explainability and stakeholder trust. This improves explanations by highlighting data relationships, understanding complex structures, and enhancing explanation quality, as demonstrated by Doctor XAI [59] and its application in decision trees [60]. It also improves interoperability in case base descriptions, promoting universal understanding and facilitating collaboration across CBR approaches.

For instance, in ADR prediction, each case can be formally represented using ontology-driven descriptors such as those defined in the literature. For example, the AI task (mentioned above in iSeeOnto) is framed as a binary classification of ADR risk using a ConvNet architecture trained on structured Electronic Health Records (EHR) data. Domain knowledge may include pharmacological interaction rules or dosage adjustment guidelines (e.g., "Avoid non-steroidal anti-inflammatory drugs, NSAIDs, in patients with impaired renal function"). These rules are not part of the learned model but are encoded as external domain knowledge used to generate adaptation explanations. This knowledge is used to adapt solutions and to explain adaptations: "The system reduced the dosage of drug X because of the patient's GFR level, following guideline G5." Explanations are presented through SHAP-based visual overlays and textual summaries designed for diverse end-users, including clinicians,

³ <https://w3id.org/iSeeOnto/explanationexperience>

pharmacists, and regulatory managers. The system enables efficient case discovery, semantic retrieval, and tailored explanation delivery across roles using a structured ontology-based case description by capturing dimensions, such as stakeholder, presentation format, intent, and technical components.

2) *RQ2.2. Generate Intuitive and Relatable Explanations Via Experience-Based Reasoning & Similarity-Based Retrieval:* Literature reveals CBR methods have been applied across various explanation types, including feature attribution, counterfactuals, and example-based explanations, showcasing CBR's versatility in creating understandable AI explanations [5], [16]. This section examines how techniques such as similarity-based retrieval [61] and experience-based reasoning [62] provide insightful explanations. Traditional CBR systems handle retrieval, matching, and adaptation in separate, sequential steps. Integrating these steps, especially retrieval and adaptability, can improve CBR system performance [63]. Gates et al.'s study [51] on case-based explanations impacts stakeholder perceptions, emphasizing solution accuracy and problem-solution similarity. Similarity-based retrieval utilizes a similarity metric to identify cases that mirror the current problem, facilitating the discovery of viable solutions. The CBR case-based approach is inherently interpretable, contrasting with complex ML techniques, and supports XAI by providing understandable explanations through experience-based reasoning [49]. Experience-based reasoning leverages knowledge from previous cases to address new challenges, drawing on insights from past experiences for informed decision-making. Combining these approaches enhances system explainability, allowing it to communicate the rationale behind solution choices and their applicability to new problems. This approach bolsters transparency and trust by showing stakeholders how decisions align with past similar cases.

Effective retrieval in XCBR systems depends on accurately determining case similarity, offering example-based explanations, also called instance-based explanations or factual explanations, based on previous instances solved by the AI model [11]. Such instance-based explanations are intuitive, matching stakeholders' mental models and learning from past experiences, making CBR one of the most effective strategies in XAI [11]. This process is essential for developing the case base and its adaptation during reuse, which will be discussed in Section IV-B3. Lamy et al. [64] introduced an explainable CBR system featuring a visual interface for breast cancer management that integrates quantitative and qualitative insights. This system merges automatic algorithmic processes with visual explanations, allowing stakeholders to visually classify queries by showing similarities with past cases. It utilizes Multidimensional Scaling for quantitative analysis and rainbow boxes for qualitative insights, offering accuracy comparable to k-Nearest Neighbors (k-NN) [65] with better explainability. Adaptation-guided retrieval (AGR) [63] streamlines case similarity assessment, minimizing the need for complex adaptation efforts. The retrieval phase often employs the MAC/FAC approach [66], a two-step process of filtering and assessing similarity. Initially, relevant cases are filtered from the case base based on the query's problem description. This filtering is direct, focusing on identifying necessary features. The subsequent similarity step,

crucial for retrieval, compares the query with filtered cases to find similarities. While CBR supports transparency in prediction understanding, the clarity of feature contributions to case ranking might be lacking [67]. Identifying similarity metrics that clarify stakeholder ranking beyond optimizing solution fit is essential. CBR's potential for enhancing understanding through visual case comparisons has also been demonstrated in the literature [61].

Selecting an appropriate similarity metric is pivotal, as it significantly impacts solutions' effectiveness and adaptability in subsequent applications. Therefore, accurately evaluating case similarity is indispensable for generating example-based explanations and implementing CBR in XAI contexts. Cunningham [19] classifies similarity metrics in CBR into four categories: direct similarity mechanisms, transformation-based measures, information-theoretic measures, and machine learning-based metrics. *Direct similarity mechanisms* utilize feature vectors to represent cases in the case base, using k-NN, and including metrics such as Overlap, euclidean, or Manhattan distances. The calculations in this group are straightforward, and stakeholders can view the similarity value or reasoning process, such as in a visualization mode. In *transformation-based measures*, similarity is considered as a transformation effort, and these measures gauge the effort required to transform one case into another. This effort is quantified as the similarity value between the two cases, and examples include Edit distances, Tree Edit Distance (TED), Levenshtein distance, or Affine Edit Distance (AED) to find the best solutions for stakeholders. Meanwhile, *information-theoretic measures* operate using the raw case representation, removing the need to create the feature vector representation. For example, information-based similarity metrics for biological sequences can reveal the differences between cases based on their raw data, directly comparing the sequence properties, such as the length of the sequence or the number of distinct elements that the sequences contain. Finally, *machine learning-based metrics* employ ML techniques, such as random forests or cluster kernels, to describe the cases. For example, we can identify the underlying properties of the instances in our dataset by applying clustering. As instances within the same cluster tend to share similar characteristics, while those in different clusters exhibit distinct ones. This insight enables meaningful comparisons between groups. An explanation method is needed to interpret and validate the ML model's behavior based on these metrics.

Similarity metrics are differentiated by their reliance on knowledge or features to compute the similarity between cases, falling into three categories: local, global, or quasi-local [68]. *Local* metrics assess similarity based on a single attribute, marking cases as similar if they share a specific property value. *Global* metrics, in contrast, evaluate all available properties and data, offering a comprehensive approach that can identify more precisely similar cases. Finally, *quasi-local* metrics utilize a selected property subset, considering similar cases if they match these chosen attributes. Local metrics may need more comprehensiveness while offering simplicity and transparency by focusing on a singular property. Global metrics, though more accurate due to their all-encompassing analysis, sacrifice transparency

as they incorporate multiple properties, increasing complexity. Furthermore, similarity metrics can incorporate structural or semantic knowledge [69]. *Structural knowledge* represents solutions using structures such as graphs or trees [18], while *semantic knowledge* uses concepts and properties defined in domain-specific ontologies [54]. Combining these allows leveraging both properties for enhanced similarity assessment, as demonstrated in [6], as part of the iSee platform, where solutions are represented by Behaviour Trees (BTs) and explained through ontology. The cloud-based CBR system, CloodCBR [70], evaluates the similarity between the context descriptions of the query case and those in the case base. It employs similarity metrics, including Wu & Palmer, Query Intersection, and Exact Match to align specific attributes or properties of a case within their iSee ontology.

In our XCBR-based ADR prediction use case, similarity retrieval is central in generating intuitive, example-based explanations. For instance, using a direct similarity metric, such as euclidean distance, the system may retrieve a case involving a patient with comparable lab values and medications, justifying the prediction with: *“This patient closely resembles a previous case with similar kidney function and drug exposure, where an ADR occurred.”* These similarity metrics also differ in their reliance on knowledge scope. A local metric might compare only renal function and medication data, offering tightly focused and interpretable results. A global metric would compute similarity across the full clinical profile, incorporating demographics, comorbidities, and lab values. A quasi-local metric, by contrast, selects a context-sensitive subset—such as SHAP-identified features or user-specified attributes—to tailor similarity and explanations dynamically: *“Similarity computed using the most influential clinical features—polypharmacy, GFR, and drug Y interaction—based on stakeholder intent and model attribution.”* Together, these similarity approaches allow XCBR systems to generate intuitive, context-aware explanations grounded in past clinical experience and adapted to the reasoning needs of diverse stakeholders.

3) *RQ2.3. Ensure Continual Relevance and Accuracy of Explanations Through Adaptation and Learning:* We will investigate the adaptation and learning features, examining how the system refines and applies insights from prior cases during the CBR reuse phase in an XCBR system. By adopting leading practices from CBR and XAI, we aim to improve case adaptation and reuse, thereby clarifying AI decision-making and ensuring the accuracy of solutions. Once the best case is selected during the retrieval phase, it may need to be modified to fit the new context. Minor tweaks may suffice for diagnostic cases, but more intricate situations, such as bespoke design challenges, demand substantial alterations for successful implementation. Considering a case base cannot encompass every possible design variation, it is vital to adapt cases to meet new demands. Transformational and generative adaptation techniques are recognized for effectiveness in this adaptation process.

Transformational adaptation [26] entails significant structural changes to a solution, modifying its fundamental composition. It involves transforming a solution from a previously retrieved case to fit a new problem until it achieves a

‘consistent’ or ‘adequate’ structure [71]. This adaptation can be detailed further, with compositional adaptation [26] representing a technique where components from various case solutions are adjusted and combined into a unified solution. Substitutional and structural adaptations are two different methods based on the degree of modification to the solution. Substitutional adaptation [26] is suited for simple problem-solving, requiring only minor changes when the retrieved case closely matches the new problem. It involves tweaking specific solution parameters or attributes without changing the core structure. Attributes from the original case are mostly kept, with adjustments made to improve the solution based on the new problem’s attributes compared to those of a similar, previously solved case. Structural adaptation [26], [63], on the other hand, significantly alters a solution’s components, such as adding or removing features, to tackle new issues beyond the scope of existing cases.

Generative adaptation or *constructive adaptation* [26], [63] is a complex adaptation form that reevaluates the reasoning process for a new problem. Unlike simple adjustments, it entails thoroughly reassessing the new context and creating new solution components rather than altering existing ones. Constructive adaptation [71], a type of generative reuse, constructs solutions for new problems by leveraging similar case solutions or parts thereof, employing a search-based approach for configuration tasks.

Adaptation is crucial during the revision phase of the CBR cycle, especially when feedback on a recommended solution indicates the need to modify a proposed solution. This phase focuses on evaluating and implementing the solution generated during the reuse phase against the stakeholder’s query, examining its fit and effectiveness while considering external system knowledge. Adjustments are made to ensure the solution meets the query’s requirements. Once an optimal solution is identified, it is vital to archive this case in the case base for future use, thereby continuously enriching the case base with new insights and enhancing the system’s problem-solving efficiency [72].

Adaptation techniques range from simple transformations to complex generative constructions. In the context of ADR prediction, these methods allow the system to adjust dosage, substitute medications, combine treatment strategies, or construct entirely new recommendations. Such strategies are particularly valuable during the revision phase, where stakeholder feedback (will discuss later) necessitate modifying the initial solution. For instance, adaptation ensures that retrieved solutions are tailored to the current patient’s context. If a retrieved case suggests a standard drug dosage, transformational adaptation may reduce the dose based on the patient’s renal function. The system explains: *“Dosage adjusted by 50% due to impaired kidney function, following renal dosing guidelines.”* In cases requiring a novel response, generative adaptation may synthesize an alternative, e.g., replacing a contraindicated drug due to a known allergy—with the explanation: *“Alternative selected from the same drug class due to allergy risk.”* Once revised and validated, the adapted case and its explanation are retained in the case base. This enables continual learning and improving the system’s ability to generate relevant, patient-specific explanations in the future.

4) *RQ2.4. Ensure Optimal Performance and Relevancy via Case Base Maintenance*: This section discusses the significance of CBM during the CBR retain phase. CBM [3] is pivotal in ensuring the effectiveness and reliability of explanations provided by the XCBR system, focusing on updating the case base to reflect current knowledge and trends for accurate and understandable explanations. Maintenance involves removing obsolete or redundant cases, consolidating cases to improve reasoning, and correcting inconsistencies [3]. For example, consolidating multiple cases of mild drug-induced rashes into a generalized dermatologic ADR case. Lupin et al. [73] describe strategies for optimizing the case base's structure or pruning it to maintain efficiency with fewer cases, potentially enhancing retrieval speed [74]. Research in the ML and CBR communities has demonstrated that decreasing the case base size effectively reduces retrieval times [75]. Various CBM algorithms have been developed, including the k-NN classifier for identifying and removing redundant or noisy cases [3], and instance reduction algorithms that cluster or evaluate instances for deletion [76]. Evaluating the updated case base's effectiveness involves traditional ML techniques, namely Hold-Out and Cross-Validation [73], comparing performance metrics against the original to ensure the CBR systems maintained or improved accuracy, sensitivity, and specificity.

The accuracy of explanations crucially depends on the quality of cases stored. High-quality cases ensure the system provides reliable solutions and trustworthy explanations. Diversity in the case base enhances the system's ability to tackle various problems, offering varied explanations. For example, retaining rare ADR examples, such as pediatric reactions to off-label drug use. However, evolving data or AI models can make case representations outdated, risking inaccurate explanations. Updating or removing such cases ensures relevance and protects stakeholders from misinformation [77]. Minimizing redundancy in the case base optimizes retrieval [78], avoiding repetitive or conflicting explanations. For instance, removing duplicate cases of elderly patients on the same drug combination with identical lab results. Moreover, the structure of the case base significantly impacts explainability, where a well-organized case base enhances the clarity of the problem-solving process and the derivation of solutions.

5) *RQ2.5. Refine Explanations & Improve System Performance Through Iterative & Interactive Process*: Despite the growing availability of XAI techniques and open-source libraries, significant work remains to advance personalization and usability in XAI and XCBR. Many current XAI methods are post-hoc and model-agnostic, meaning they can be applied across different AI models but often lack stakeholder-specific adaptation. As a result, they may produce technically accurate explanations but not intuitive or meaningful for non-expert stakeholders. This limitation also extends to XCBR. While case-based explanations are inherently more interpretable, grounded in examples from past experience, they are frequently presented as generic instances rather than tailored explanations. In many cases, they do not highlight which retrieved cases are most relevant or understandable for a specific stakeholder, nor do they

adapt the explanation format or content based on stakeholder expertise, intent, or context.

Building on the need for greater personalization and usability in XCBR systems, exploring techniques that support stakeholder interaction and feedback becomes essential. Interaction is critical in refining explanations—aligning them more closely with the stakeholder's mental model and enabling the system to learn from stakeholder preferences, corrections, and contextual cues. Incorporating stakeholder feedback allows XCBR systems to go beyond static, example-based explanations and evolve into interactive, adaptive systems that continuously improve their relevance and interpretability. This section reviews methods that facilitate this bidirectional exchange.

Robust stakeholder interaction, feedback mechanisms, and iterative learning are crucial for maintaining a stakeholder-centric approach and fostering continual improvement in XCBR systems [50], [79]. Enabling stakeholders to shape and tailor explanations enhances the transparency of a predictive system, making it more relevant and engaging to the explainee or recipient by aligning the explanation's detail and complexity with their preferences. Hence, personalization can be understood as the adaptation of an explanation or the explanatory process to address the unique queries of individual stakeholders [50]. For instance, a clinician may receive an alert indicating that a patient is at high risk for an adverse drug reaction. While a generic system might state: "Risk flagged due to polypharmacy and impaired renal function," a personalized and interactive system would allow the clinician to ask, "How could this risk have been reduced?" The system could generate a counterfactual explanation in response, such as "If the patient were not concurrently prescribed drug X, the predicted ADR risk would fall below the threshold." This clarifies the rationale behind the alert and supports clinical decision-making by suggesting modifiable factors that could change the outcome.

Developing user-friendly and intuitive interfaces and aligning with stakeholders' perspectives is crucial to enhancing system interaction. These interfaces need to evolve based on stakeholder behavior and feedback, aiming for a more intuitive and stakeholder-centered design. Incorporating simple feedback mechanisms, such as ratings, comments, or interactive dialogues, facilitates open communication, allowing stakeholders to share their views and suggestions quickly. Matrix factorization-based collaborative filtering, for instance, utilizes user ratings to find similarities among stakeholders or items, creating personalized recommendations [80]. Despite its effectiveness, matrix factorization is often seen as a black box. Addressing this, a post-hoc, model-agnostic explanation system for matrix factorization recommendations [81] is introduced, employing CBR and Formal Concept Analysis (FCA). This system creates a case base of items linked to recommended items by analyzing stakeholder interactions, identifying explanatory items through cosine similarity, and using FCA to identify key feature groupings.

The posed questions mold causal explanations in conversations and follow specific dialogue principles [82]. A causal conversation model posits that adequate explanations must be

accurate and directly address the core of the ‘why’ question. This relevance principle links various attribution models through common counterfactual reasoning tailored to different causal questions. For instance, a clinician inquires, “Why is this patient flagged as high risk?” The system could respond with a structured explanation: “The high ADR risk is due to the concurrent use of drug X and reduced renal function.” An explanatory agent within the XCBR system could then elaborate using multiple forms of communication presenting a textual summary, highlighting the key contributing features in a visual plot, and offering an interactive comparison with similar past cases. This illustrates how an explanatory agent could function by answering questions and using multiple communication forms, including text and visuals, to clarify its decision-making process [13]. In human-AI interactions [83], dialogue-like explanations with visual enhancements are most impactful. This method improves AI transparency, making it more adaptable and understandable to the stakeholder, and allows for a deeper exploration of fairness and accountability [13]. It also promotes bilateral communication, letting stakeholders contribute their insights to refine the AI’s logic.

Surveys and feedback forms are crucial for collecting stakeholder feedback on system operations, explanations, and decision-making. Smith et al. [79] found that providing explanations without feedback opportunities, especially with lower-quality models, frustrates stakeholders. Conversely, high-quality models require detailed feedback with explanations. Streamlining the feedback process boosts stakeholder engagement and values their input. A graph-based approach [84] leverages link prediction to tailor explanations based on stakeholder activity, enhancing personalization. Feedback-driven updates to the case base are critical, with stakeholder feedback guiding case revisions. Positive feedback confirms case effectiveness, while negative feedback highlights revision needs. Nick et al. [85] proposed an experience feedback loop for CBM driven by stakeholder feedback, addressing case acquisition and maintenance challenges. Liao et al. [86] introduced a Question-Driven Design Process focused on user questions to select and evaluate XAI techniques, enhancing the stakeholder experience. Adapting XAI algorithms based on feedback improves system reasoning and decision-making [87]. An iterative case refinement process involving case description and solution adjustments based on stakeholder input fills identified gaps. Sosa et al. [88] developed IREX for Iterative Refinement and EXplanation of classification models, allowing domain experts to refine models for tabular data through iterative expert feedback and explanation techniques, optimizing model performance. Adapting similarity metrics based on stakeholder feedback can increase case retrieval relevance [89]. If stakeholders prefer specific case types, adjusting the weights of those features can enhance retrieval accuracy.

Implementing a ‘human-in-the-loop’ strategy [4] achieves equilibrium between the capabilities of AI and human supervision, enabling stakeholders to identify and report biases and impacts. Encouraging stakeholder contributions can diversify the case base with various real-world examples. The system

should promptly respond to feedback for improved interaction and stakeholder experience.

6) *RQ2.6. Evaluate Performance and Explainability Through Strategies & Metrics:* This section focuses on the evaluation metrics for XCBR. Performance measurement aims to evaluate how well the human-machine system performs the tasks it was created to handle. Hoffman et al. [14] suggest that providing stakeholders with satisfactory explanations should enhance the overall effectiveness of the stakeholder and the AI system. Additionally, evaluating an XAI system’s effectiveness is intertwined with the stakeholder’s performance and that of the human-machine system. Research on AI/ML model explainability diverges into two main paths: model complexity and human-centric evaluations [15]. Model complexity evaluations associate interpretability with simplicity, using criteria such as model size and rule count [90], with studies, including scalable Bayesian Rule Lists focusing on achieving a balance between interpretability, accuracy, and efficiency [90]. Recent efforts aim to quantify model complexity through model-agnostic metrics, observing how feature interactions affect explainability, such as in ALE Plots [91]. Meanwhile, human-centric studies evaluate explainability through formats, such as decision tables and trees, exploring their effects on interpretation and trust [92], and investigate the influence of transparency and feature quantity on stakeholder trust and prediction accuracy [93]. Evaluations of XAI systems incorporate stakeholder feedback and quantitative analyses to gauge their impact [94].

Online and offline evaluations are two primary methods to assess the effectiveness and quality of explanations in XCBR systems. Online evaluations involve direct stakeholder interaction for system assessment, which is crucial for understanding stakeholder experience, satisfaction, and the practical value of explanations. Tools for online evaluation often include surveys and questionnaires utilizing formats, including Likert scales, open-ended questions, or targeted inquiries about explanation clarity, relevance, and utility [81]. Researchers also use metrics, namely Net Promoter Scores [95] or user satisfaction indices, to measure general acceptance and perceived explanation quality. Comparing stakeholder perceptions before and after interacting with the system highlights the explanations’ concrete effects.

Offline evaluation in XCBR systems leverages pre-defined metrics to measure explanation quality objectively without relying on stakeholder feedback. This approach offers structured, repeatable assessments and ensures unbiased evaluation. Fidelity assessments [96] gauge the accuracy of explanations in mirroring the model’s internal workings, aiming for a precise match between the model’s actions and the explanations provided. Consistency checks ensure that similar cases yield similar explanations, indicating the process’s reliability. Coverage and diversity metrics [97] evaluate the system’s capacity to explain various cases and the range of explanations generated, highlighting the system’s comprehensive explanatory capability and flexibility. Correlation assessments examine how closely explanations align with the system’s reasoning, typically validated by comparing outcomes to known correct responses or expert

evaluations. Quantitative measures such as precision, recall, and F1-score are employed to quantitatively validate explanation quality, offering a clear metric for assessing the effectiveness of explanations in XCBR systems. Recent studies introduce offline metrics for analyzing various attributes of explanations, including dissimilarity, sparsity, instability, or runtime in counterfactual explanations [45]. This shift underscores the ongoing need to develop evaluation metrics for XAI and highlights the current focus within the XCBR domain.

V. CHALLENGES, OPPORTUNITIES, AND FUTURE DIRECTIONS IN XCBR

This section tackles RQ3, outlined in Section II, focusing on advancing AI to better meet human needs through interpretability, fairness, and accountability. Integrating CBR and XAI can improve transparency in decision-support systems; however, various challenges still hinder XCBR's effectiveness. A major challenge is the complexity of case retrieval, which requires domain-specific similarity metrics responsive to the data's quality, dimensionality, and completeness. For example, high-dimensional clinical datasets face challenges related to missing features and diverse information sources, which affect retrieval effectiveness. Another major challenge arises from the object-oriented nature of CBR explanations [2]. Unlike many standard AI systems, XCBR must address the needs of diverse stakeholders beyond end-users based on their roles, expertise, and information needs [2]. XCBR faces significant difficulties in managing the diversity of interpretations required for heterogeneous case bases [2]. As case libraries contain objects with varying structures, semantics, and feature spaces, constructing explanation models that generalize across this diversity remains an open challenge. Technical challenges are heightened by data limitations, including small sample sizes, class imbalance, label noise, and inadequate representation of rare conditions, all hindering system generalizability. As case bases grow, the demand for scalable indexing and practical similarity search algorithms rises to support prompt decision-making in critical areas, particularly healthcare. Moreover, maintaining consistent and reliable explanations during case base expansion poses additional challenges for long-term system maintenance [2].

XAI is essential for advancing responsible AI [15], [98] while XCBR faces challenges in developing trustworthy systems. A universal XCBR framework is needed to support diverse stakeholders by standardizing case retrieval, adaptation, and explanation mechanisms while ensuring explanations remain understandable and meaningful across clinical, regulatory, and technical domains. Establishing such a framework would allow systematic evaluation of explanation effectiveness, including stakeholder trust, comprehension, and clinical utility, through carefully designed human studies. While CBR offers an intuitive foundation for generating explanations through prior cases, its broader adoption in XAI remains limited, particularly when dealing with highly complex or high-dimensional AI models. Current CBR contributions to XAI have demonstrated value in domains such as image classification, text analysis, and neural network interpretation; however, the full potential of XCBR for

complex decision support tasks has yet to be fully realized [2]. The hybrid integration of CBR with model-based XAI techniques, such as SHAP, Anchors, and Influence Functions, offers a pathway toward richer, multi-level explanations that combine the strengths of precedent-based reasoning with data-driven model interpretability. In addition, improving human-AI collaboration [4] within XCBR requires the development of adaptive explanation strategies that can modify explanations according to the stakeholder's level of expertise, information needs, and contextual elements. Drawing insights from philosophy, psychology, and cognitive science can refine XCBR systems to align with human explanation processing and evaluation [4]. Moreover, incorporating knowledge from various sectors can further develop human-centric XCBR approaches.

The evolving landscape of XAI increasingly emphasizes auditing AI models and enhancing the usability and comprehensibility of XAI interfaces. While many approaches offer explanations for complex models, they often lack customization for non-expert stakeholders. XCBR provides explanations based on past cases but may not effectively identify the most relevant examples for specific stakeholders. Future research should focus on developing interactive XCBR tools that adapt to various stakeholders' needs and thought processes. This transition from static to interactive explanations is essential for the future of explainable, stakeholder-centered AI. Key focus areas should include personalized systems incorporating stakeholder models, feedback mechanisms, and context sensitivity to improve interpretability, trustworthiness, and decision-making support.

Development challenges in responsible human-centered AI deployment focus on enhancing communication of data quality and addressing analysis sparsity through innovative methods [98]. Future advancements in XCBR require new design patterns, knowledge graphs, ontology creation methodologies [98], [99], and integration techniques with AI. User studies are crucial for evaluating how ontologies enhance understanding and transparency. Effective communication of data quality to stakeholders, addressing completeness and fairness, is vital for responsible AI development. Another probable direction for advancement is the integration of Large Language Models (LLMs) [100] into XCBR systems. The potential in these models lies in improving the ability to generate clear and natural language explanations by making the systems more engaging and user-friendly to interact with. This aligns well with ongoing AI research, where efforts are being directed towards creating precise, transparent systems centered on the stakeholder's experience [101]. The future of XCBR involves developing frameworks that clarify AI processes, meet diverse stakeholder needs, uphold ethical principles, and support transparent human-AI interactions, necessitating a multidisciplinary approach.

VI. CONCLUSION

In this paper, we examined XCBR through a cycle-aligned perspective that examined how explainability can be systematically supported across the four core phases of the CBR process: Retrieve, Reuse, Revise, and Retain. By merging the strengths of CBR, which offers inherently interpretable,

stakeholder-focused, case-based explanations, with XAI for transparency and accountability, we positioned XCBR as a powerful intersection of these paradigms. XCBR systems utilizes CBR for case-based explanations or generating explanations using CBR techniques, a significant advancement in demystifying decision-making. Our work identified six functional goals for explainability, including semantic interpretation, intuitive & relatable explanations, continual relevance & accuracy of explanations, optimal performance & relevance, refine explanations, and evaluate the performance and explainability. These goals were mapped to six thematic categories: structured case representation, domain knowledge, experience-based reasoning, similarity retrieval, case adaptation, case storage and management, user feedback, and performance evaluation, forming the foundation of a comprehensive XCBR taxonomy. This taxonomy serves a descriptive framework and a diagnostic tool to analyze, compare, and design XCBR systems.

There remain challenges in building interactive, adaptive, and scalable XCBR systems that align with diverse stakeholder needs and real-world constraints despite CBR's natural compatibility with XAI objectives. There are also opportunities for enhancing stakeholder-personalized explanations, employing LLMs to generate explanations, and integrating human-centric measures into evaluation procedures. In this paper, we review the significance of ongoing research and innovation in XCBR to fulfill the evolving requirements of AI deployment. Continued collaboration, methodological care, and collective effort are needed to develop genuinely explainable and trustworthy AI to build systems that benefit humanity.

REFERENCES

- [1] C. Zednik, "Solving the black box problem: A normative framework for explainable artificial intelligence," *Philosophy Technol.*, vol. 34, no. 2, pp. 265–288, 2021.
- [2] A. Yan and Z. Cheng, "A review of the development and future challenges of case-based reasoning," *Appl. Sci.*, vol. 14, no. 16, 2024, Art. no. 7130.
- [3] B. Chebel-Morello, M. K. Haouchine, and N. Zerhouni, "Case-based maintenance: Structuring and incrementing the case base," *Knowl.-Based Syst.*, vol. 88, pp. 165–183, 2015.
- [4] V. Hassija et al., "Interpreting black-box models: A review on explainable artificial intelligence," *Cogn. Computation*, vol. 16, no. 1, pp. 45–74, 2024.
- [5] M. T. Keane and E. M. Kenny, "How case-based reasoning explains neural networks: A theoretical analysis of xai using post-hoc explanation-by-example from a survey of ANN-CBR twin-systems," in *Proc. Case-Based Reasoning Res. Development: 27th Int. Conf.*, Otzenhausen, Germany, 2019, pp. 155–171.
- [6] A. Wijekoon et al., "CBR driven interactive explainable AI," in *Proc. Int. Conf. Case-Based Reasoning*, 2023, pp. 169–184.
- [7] J. M. Schoenborn, R. O. Weber, D. W. Aha, J. Cassens, and K.-D. Althoff, "Explainable case-based reasoning: A survey," in *Proc. AAAI-21 Workshop*, 2021.
- [8] L. Gates and D. Leake, "Evaluating CBR explanation capabilities: Survey and next steps," in *Proc. ICCBR Workshops*, 2021, pp. 40–51.
- [9] R. Weber, M. Shrestha, and A. J. Johs, "Knowledge-based XAI through CBR: There is more to explanations than models can tell," 2021, *arXiv:2108.10363*.
- [10] S. R. Islam and W. Eberle, "Domain knowledge-aided explainable artificial intelligence," in *Proc. Explainable Artif. Intell. Cyber Secur.: Next Gener. Artif. Intell.*, 2022, pp. 73–92.
- [11] J. van der Waa, E. Nieuwburg, A. Cremers, and M. Neerincx, "Evaluating xai: A comparison of rule-based and example-based explanations," *Artif. Intell.*, vol. 291, 2021, Art. no. 103404.
- [12] A. Adadi and M. Berrada, "Peeking inside the black-box: A survey on explainable artificial intelligence (XAI)," *IEEE Access*, vol. 6, pp. 52138–52160, 2018.
- [13] T. Miller, "Explanation in artificial intelligence: Insights from the social sciences," *Artif. Intell.*, vol. 267, pp. 1–38, 2019.
- [14] R. R. Hoffman, S. T. Mueller, G. Klein, and J. Litman, "Metrics for explainable AI: Challenges and prospects," 2018, *arXiv: 1812.04608*.
- [15] S. R. Islam, W. Eberle, S. K. Ghafoor, and M. Ahmed, "Explainable artificial intelligence approaches: A survey," 2021, *arXiv:2101.09429*.
- [16] A. Aamodt and E. Plaza, "Case-based reasoning: Foundational issues, methodological variations, and system approaches," *AI Commun.*, vol. 7, no. 1, pp. 39–59, 1994.
- [17] S. H. El-Sappagh and M. Elmogy, "Case based reasoning: Case representation methodologies," *Int. J. Adv. Comput. Sci. Appl.*, vol. 6, no. 11, pp. 192–208, 2015.
- [18] H. Xu, Y. Wei, Y. Cai, and B. Xing, "Knowledge graph and cbr-based approach for automated analysis of bridge operational accidents: Case representation and retrieval," *PLoS One*, vol. 18, no. 11, 2023, Art. no. e0294130.
- [19] P. Cunningham, "A taxonomy of similarity mechanisms for case-based reasoning," *IEEE Trans. Knowl. Data Eng.*, vol. 21, no. 11, pp. 1532–1543, Nov. 2009.
- [20] M. Caro-Martínez, A. Wijekoon, J. A. Recio-García, D. Corsar, and I. Nkisi-Orji, "Conceptual modelling of explanation experiences through the iSeeOnto ontology," in *Proc. CEUR Workshop*, 2023.
- [21] M. Caro-Martínez, G. Jiménez-Díaz, and J. A. Recio-García, "Conceptual modeling of explainable recommender systems: An ontological formalization to guide their design and development," *J. Artif. Intell. Res.*, vol. 71, pp. 557–589, 2021.
- [22] F. Liqing and A. Senthil Kumar, "XML-based representation in a CBR system for fixture design," *Comput.-Aided Des. Appl.*, vol. 2, no. 1/4, pp. 339–348, 2005.
- [23] A. B. Arrieta et al., "Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI," *Inf. Fusion*, vol. 58, pp. 82–115, 2020.
- [24] R. L. De Mantaras et al., "Retrieval, reuse, revision and retention in case-based reasoning," *Knowl. Eng. Rev.*, vol. 20, no. 3, pp. 215–240, 2005.
- [25] M. M. Richter and M. Michael, "Knowledge containers," in *Readings in Case-Based Reasoning*, San Mateo, CA, USA: Morgan Kaufmann, 2003.
- [26] W. Wilke and R. Bergmann, "Techniques and knowledge used for adaptation during case-based problem solving," in *Proc. Tasks Methods Appl. Artif. Intell.: 11th Int. Conf. Ind. Eng. Appl. Artif. Intell. Expert Syst. IEA-98-AIE Benicàssim*, Castellón, Spain, 1998, pp. 497–506.
- [27] K. Martin, "Similarity and explanation for dynamic telecommunication engineer support," PhD dissertation, 2021, doi: [10.48526/rgu-wt-1447160](https://doi.org/10.48526/rgu-wt-1447160).
- [28] A. F. Markus, J. A. Kors, and P. R. Rijnbeek, "The role of explainability in creating trustworthy artificial intelligence for health care: A comprehensive survey of the terminology, design choices, and evaluation strategies," *J. Biomed. Informat.*, vol. 113, 2021, Art. no. 103655.
- [29] B. Goodman and S. Flaxman, "European union regulations on algorithmic decision-making and a 'right to explanation'," *AI Mag.*, vol. 38, no. 3, pp. 50–57, 2017.
- [30] M. Langer et al., "What do we want from explainable artificial intelligence (XAI)?—a stakeholder perspective on XAI and a conceptual model guiding interdisciplinary XAI research," *Artif. Intell.*, vol. 296, 2021, Art. no. 103473.
- [31] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and D. Pedreschi, "A survey of methods for explaining black box models," *ACM Comput. Surv.*, vol. 51, no. 5, pp. 1–42, 2018.
- [32] I. Stepin, J. M. Alonso, A. Catala, and M. Pereira-Fariña, "A survey of contrastive and counterfactual explanation generation methods for explainable artificial intelligence," *IEEE Access*, vol. 9, pp. 11974–12001, 2021.
- [33] M. T. Keane and B. Smyth, "Good counterfactuals and where to find them: A case-based technique for generating counterfactuals for explainable AI (XAI)," in *Proc. Case-Based Reasoning Res. Development: 28th Int. Conf.*, Salamanca, Spain, 2020, pp. 163–178.
- [34] W. V. Woensel et al., "Explainable clinical decision support: Towards patient-facing explanations for education and long-term behavior change," in *Proc. Int. Conf. Artif. Intell. Med.*, 2022, pp. 57–62.
- [35] R. Dwivedi et al., "Explainable AI (XAI): Core ideas, techniques, and solutions," *ACM Comput. Surv.*, vol. 55, no. 9, pp. 1–33, 2023.

- [36] S. Sawangreerak, P. Thanathamath, P. Lakkanawant, and N. S. A. Wahab, "Anchor-based explainable and causal artificial intelligence for enhancing financial predictions of future earnings," *IEEE Access*, to be published, doi: [10.1109/ACCESS.2025.3557264](https://doi.org/10.1109/ACCESS.2025.3557264).
- [37] T. N. Mundhenk, B. Y. Chen, and G. Friedland, "Efficient saliency maps for explainable AI," 2019, *arXiv: 1911.11293*.
- [38] R. Guidotti, A. Monreale, F. Giannotti, D. Pedreschi, S. Ruggieri, and F. Turini, "Factual and counterfactual explanations for black box decision making," *IEEE Intell. Syst.*, vol. 34, no. 6, pp. 14–23, Dec. 2019.
- [39] S. Aryal and M. T. Keane, "Even if explanations: Prior work, desiderata & benchmarks for semi-factual XAI," 2023, *arXiv:2301.11970*.
- [40] P. W. Koh and P. Liang, "Understanding black-box predictions via influence functions," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 1885–1894.
- [41] R. Grosse et al., "Studying large language model generalization with influence functions," 2023, *arXiv:2308.03296*.
- [42] V. Arya et al., "AI explainability 360 toolkit," in *Proc. 3rd ACM India Joint Int. Conf. Data Sci. Manage. Data (8th ACM IKDD CODS 26th COMAD)*, 2021, pp. 376–379.
- [43] N. Kokhlikyan et al., "Captum: A unified and generic model interpretability library for PyTorch," 2020, *arXiv: 2009.07896*.
- [44] A. Wijekoon and N. Wiratunga, "A user-centred evaluation of discern: Discovering counterfactuals for code vulnerability detection and correction," *Knowl.-Based Syst.*, vol. 278, 2023, Art. no. 110830.
- [45] B. Bayrak and K. Bach, "PerTCF: A perturbation-based counterfactual generation approach," in *Proc. Int. Conf. Innov. Techn. Appl. Artif. Intell.*, 2023, pp. 174–187.
- [46] A. B. Tofighi, A. Ahmadi, and H. Mosadegh, "A novel case-based reasoning system for explainable lung cancer diagnosis," *Comput. Biol. Med.*, vol. 185, 2025, Art. no. 109547.
- [47] P. Pradeep, M. Caro-Martínez, and A. Wijekoon, "A practical exploration of the convergence of case-based reasoning and explainable artificial intelligence," *Expert Syst. Appl.*, vol. 255, 2024, Art. no. 124733.
- [48] R. Bergmann, G. Pews, and W. Wilke, "Explanation-based similarity: A unifying approach for integrating domain knowledge into case-based reasoning for diagnosis and planning tasks," in *Proc. Topics Case-Based Reasoning: First Eur. Workshop*, Germany, 1994, pp. 182–196.
- [49] A. J. Cañas, D. B. Leake, and A. G. Maguitman, "Combining concept mapping with CBR: Towards experience-based support for knowledge modeling," in *Proc. FLAIRS Conf.*, Citeseer, 2001, pp. 286–290.
- [50] K. Sokol and P. Flach, "One explanation does not fit all: The promise of interactive explanations for machine learning transparency," *KI-Künstliche Intelligenz*, vol. 34, no. 2, pp. 235–250, 2020.
- [51] L. Gates, D. Leake, and K. Wilkerson, "Cases are king: A user study of case presentation to explain CBR decisions," in *Proc. Int. Conf. Case-Based Reasoning*, 2023, pp. 153–168.
- [52] Y. Guo, B. Zhang, Y. Sun, K. Jiang, and K. Wu, "Machine learning based feature selection and knowledge reasoning for CBR system under Big Data," *Pattern Recognit.*, vol. 112, 2021, Art. no. 107805.
- [53] P. Pradeep and S. Krishnamoorthy, "The MOM of context-aware systems: A survey," *Comput. Commun.*, vol. 137, pp. 44–69, 2019.
- [54] A. Abou Assali, D. Lenne, and B. Debray, "Case retrieval in ontology-based CBR systems," in *Proc. KI 2009: Adv. Artif. Intell.: 32nd Annu. German Conf. AI*, Paderborn, Germany, 2009, pp. 564–571.
- [55] M. Kulmanov, F. Z. Smaili, X. Gao, and R. Hoehndorf, "Semantic similarity and machine learning with ontologies," *Brief. Bioinf.*, vol. 22, no. 4, 2021, Art. no. bbaa 199.
- [56] S. Chari, O. Seneviratne, D. M. Gruen, M. A. Foreman, A. K. Das, and D. L. McGuinness, "Explanation ontology: A model of explanations for user-centered AI," in *Proc. Int. Semantic Web Conf.*, 2020, pp. 228–243.
- [57] I. Tiddi, M. d'Aquin, and E. Motta, "An ontology design pattern to define explanations," in *Proc. 8th Int. Conf. Knowl. Capture*, 2015, pp. 1–8.
- [58] I. Padhiar, O. Seneviratne, S. Chari, D. Gruen, and D. L. McGuinness, "Semantic modeling for food recommendation explanations," in *Proc. IEEE 37th Int. Conf. Data Eng. Workshops*, 2021, pp. 13–19.
- [59] C. Panigutti, A. Perotti, and D. Pedreschi, "Doctor XAI: An ontology-based approach to black-box sequential data classification explanations," in *Proc. 2020 Conf. Fairness, Accountability, Transparency*, 2020, pp. 629–639.
- [60] R. Confalonieri, T. Weyde, T. R. Besold, and F. M. del Prado Martíu, "Using ontologies to enhance human understandability of global post-hoc explanations of black-box models," *Artif. Intell.*, vol. 296, 2021, Art. no. 103471.
- [61] P. Marín-Veites and K. Bach, "Explaining CBR systems through retrieval and similarity measure visualizations: A case study," in *Proc. Int. Conf. Case-Based Reasoning*, 2022, pp. 111–124.
- [62] D. Wang, Q. Yang, A. Abdul, and B. Y. Lim, "Designing theory-driven user-centric explainable AI," in *Proc. 2019 CHI Conf. Hum. Factors Comput. Syst.*, 2019, pp. 1–15.
- [63] B. Smyth and M. T. Keane, "Using adaptation knowledge to retrieve and adapt design cases," *Knowl.-Based Syst.*, vol. 9, no. 2, pp. 127–135, 1996.
- [64] J.-B. Lamy, B. Sekar, G. Guezennec, J. Bouaud, and B. Séroussi, "Explainable artificial intelligence for breast cancer: A visual case-based reasoning approach," *Artif. Intell. Med.*, vol. 94, pp. 42–53, 2019.
- [65] T. Cover and P. Hart, "Nearest neighbor pattern classification," *IEEE Trans. Inf. Theory*, vol. 13, no. 1, pp. 21–27, Jan. 1967.
- [66] K. D. Forbus, D. Gentner, and K. Law, "MAC/FAC: A model of similarity-based retrieval," *Cogn. Sci.*, vol. 19, no. 2, pp. 141–205, 1995.
- [67] K. Bach and P. J. Mork, "On the explanation of similarity for developing and deploying CBR systems," in *Proc. 33rd Int. Flairs Conf.*, 2020, pp. 413–416.
- [68] L. Lü and T. Zhou, "Link prediction in complex networks: A survey," *Phys. A: Statist. Mech. Appl.*, vol. 390, no. 6, pp. 1150–1170, 2011.
- [69] A. Günay and P. Yolum, "Structural and semantic similarity metrics for web service matchmaking," in *Proc. Int. Conf. Electron. Commerce Web Technol.*, 2007, pp. 129–138.
- [70] I. Nkisi-Orji, N. Wiratunga, C. Palihawadana, J. A. Recio-García, and D. Corsar, "Cloud CBR: Towards microservices oriented case-based reasoning," in *Proc. Int. Conf. Case-Based Reasoning*, 2020, pp. 129–143.
- [71] E. Plaza and J.-L. Arcos, "Constructive adaptation," in *Proc. Eur. Conf. Case-Based Reasoning*, 2002, pp. 306–320.
- [72] F. Fdez-Riverola, J. M. Corchado, and J. M. Torres, "An automated hybrid CBR system for forecasting," in *Proc. Eur. Conf. Case-Based Reasoning*, 2002, pp. 519–533.
- [73] E. Lupiani, J. M. Juárez, and J. Palma, "Evaluating case-base maintenance algorithms," *Knowl.-Based Syst.*, vol. 67, pp. 180–194, 2014.
- [74] B. Smyth and E. McKenna, "Competence guided incremental footprint-based retrieval," *Knowl.-Based Syst.*, vol. 14, no. 3–4, pp. 155–161, 2001.
- [75] L. Cummins and D. Bridge, "Choosing a case base maintenance algorithm using a meta-case base," in *Proc. Int. Conf. Innov. Techn. Appl. Artif. Intell.*, 2011, pp. 167–180.
- [76] B.-R. Dai and S.-M. Hsu, "An instance selection algorithm based on reverse nearest neighbor," in *Proc. Adv. Knowl. Discov. Data Mining: 15th Pacific-Asia Conf.*, Shenzhen, China, 2011, pp. 1–12.
- [77] K. Göbel, C. Niessen, S. Seufert, and U. Schmid, "Explanatory machine learning for justified trust in human-ai collaboration: Experiments on file deletion recommendations," *Front. Artif. Intell.*, vol. 5, 2022, Art. no. 919534.
- [78] E. C. Tsang and X. Wang, "An approach to case-based maintenance: Selecting representative cases," *Int. J. Pattern Recognit. Artif. Intell.*, vol. 19, no. 1, pp. 79–89, 2005.
- [79] A. Smith-Renner et al., "No explainability without accountability: An empirical study of explanations and feedback in interactive ML," in *Proc. 2020 CHI Conf. Hum. Factors Comput. Syst.*, 2020, pp. 1–13.
- [80] G. Adomavicius and A. Tuzhilin, "Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions," *IEEE Trans. Knowl. Data Eng.*, vol. 17, no. 6, pp. 734–749, Jun. 2005.
- [81] J. L. Jorro-Aragoneses, M. Caro-Martínez, B. Díaz-Agudo, and J. A. Recio-García, "A user-centric evaluation to generate case-based explanations using formal concept analysis," in *Proc. Int. Conf. Case-Based Reasoning*, 2020, pp. 195–210.
- [82] D. J. Hilton, "Conversational processes and causal explanation," *Psychol. Bull.*, vol. 107, no. 1, 1990, Art. no. 65.
- [83] S. Amershi et al., "Guidelines for human-AI interaction," in *Proc. 2019 CHI Conf. Hum. Factors Comput. Syst.*, 2019, pp. 1–13.
- [84] M. Caro-Martínez, J. A. Recio-García, and G. Jiménez-Díaz, "An algorithm independent case-based explanation approach for recommender systems using interaction graphs," in *Proc. Case-Based Reasoning Res. Development: 27th Int. Conf.*, Ottenhausen, Germany, 2019, pp. 17–32.
- [85] M. Nick, "Reducing the case acquisition and maintenance bottleneck with user-feedback-driven case base maintenance," in *Proc. FLAIRS Conf.*, 2006, pp. 376–382.
- [86] Q. V. Liao, M. Pribić, J. Han, S. Miller, and D. Sow, "Question-driven design process for explainable ai user experiences," 2021, *arXiv:2104.03483*.
- [87] Y. Ramon, T. Vermeire, O. Toubia, D. Martens, and T. Evgeniou, "Understanding consumer preferences for explanations generated by XAI algorithms," 2021, *arXiv:2107.02624*.
- [88] C. E. Sosa-Espadas, M. G. Orozco-del Castillo, N. Cuevas-Cuevas, and J. A. Recio-García, "IREX: Iterative refinement and explanation of classification models for tabular datasets," *SoftwareX*, vol. 23, 2023, Art. no. 101420.

- [89] I. Sembiring, P. A. Christianto, and E. Sedyono, "A note on the combination of the new similarity formula with feedback to better handle complaints of in Vitro fertilization (IVF) patients," *Karbala Int. J. Modern Sci.*, vol. 9, no. 3, 2023, Art. no. 5.
- [90] H. Yang, C. Rudin, and M. Seltzer, "Scalable Bayesian rule lists," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 3921–3930.
- [91] C. Molnar, G. Casalicchio, and B. Bischl, "Quantifying model complexity via functional decomposition for better post-hoc interpretability," in *Proc. Mach. Learn. Knowl. Discov. Databases: Int. Workshops ECML PKDD 2019*, Würzburg, Germany, 2020, pp. 193–204.
- [92] J. Huysmans, K. Dejaeger, C. Mues, J. Vanthienen, and B. Baesens, "An empirical evaluation of the comprehensibility of decision table, tree and rule based predictive models," *Decis. Support Syst.*, vol. 51, no. 1, pp. 141–154, 2011.
- [93] F. Poursabzi-Sangdeh, D. G. Goldstein, J. M. Hofman, J. W. Wortman Vaughan, and H. Wallach, "Manipulating and measuring model interpretability," in *Proc. 2021 CHI Conf. Hum. Factors Comput. Syst.*, 2021, pp. 1–52.
- [94] B. Bayrak, "Beyond post-hoc instance-based explanation methods," in *Proc. ICCBR 2023 Workshop Proc.*, 2023.
- [95] I. Rallis et al., "Interpretation of net promoter score attributes using explainable AI," in *Proc. 15th Int. Conf. Pervasive Technol. Related to Assistive Environ.*, 2022, pp. 113–117.
- [96] E. M. Kenny and M. T. Keane, "Explaining deep learning using examples: Optimal feature weighting methods for twin systems using post-hoc, explanation-by-example in XAI," *Knowl.-Based Syst.*, vol. 233, 2021, Art. no. 107530.
- [97] M. Kaminskas and D. Bridge, "Diversity, serendipity, novelty, and coverage: A survey and empirical analysis of beyond-accuracy objectives in recommender systems," *ACM Trans. Interactive Intell. Syst.*, vol. 7, no. 1, pp. 1–42, 2016.
- [98] W. Saeed and C. Omlin, "Explainable AI (XAI): A systematic meta-survey of current challenges and future opportunities," *Knowl.-Based Syst.*, vol. 263, 2023, Art. no. 110273.
- [99] N. Burkart and M. F. Huber, "A survey on the explainability of supervised machine learning," *J. Artif. Intell. Res.*, vol. 70, pp. 245–317, 2021.
- [100] A. Zytek, S. Pidò, and K. Veeramachaneni, "LLMs for XAI: Future directions for explaining explanations," 2024, *arXiv:2405.06064*.
- [101] K. Wilkerson and D. Leake, "On implementing case-based reasoning with large language models," in *Proc. Int. Conf. Case-Based Reasoning*, 2024, pp. 404–417.



Preeja Pradeep received the MSc degree in computer science from Amrita Vishwa Vidyapeetham, India, in 2009, the MTech degree in wireless networks, in 2012, and the PhD degree in the area of context-aware Internet of Things (IoT) ecosystems, in 2021. She is a postdoctoral researcher working in the nasc research group, School of Computer Science and IT, University College Cork, Ireland. She has 10+ years of research experience focusing on resource scheduling and optimization, XAI, context-aware intelligent systems, IoT, data analysis, and WSN.



Marta Caro-Martínez received the PhD degree in artificial intelligence and computer science from Same University, in 2022. She is a researcher and lecturer with the Complutense University of Madrid, Spain. With eight years of experience in research, her work focuses on Artificial Intelligence, specifically on XAI, recommender systems, and case-based reasoning. She is currently the Principal investigator of the SHOU-X project, which aims to find synergies between XAI and HCI improving the usability of XAI techniques.



Anjana Wijekoon received the BSc Eng. (Hons) in computer science and engineering from the University of Moratuwa, Sri Lanka, the MSc degree in computer science from the University of Moratuwa, Sri Lanka. She is a senior research fellow working in Surgical AI with the UCL Hawkes Institute, University College London, U.K.. Previously, she was working on Explainable AI with the chist-era funded iSee project. Her PhD (2021) investigated multi-modal fusion and meta-learning approaches for personalised activity recognition with applications in the digital health domain. Prior to RGU.