

CBR driven Interactive Explainable AI*

Anjana Wijekoon¹, Nirmalie Wiratunga¹, Kyle Martin¹, David Corsar¹,
Ikechukwu Nkisi-Orji¹, Chamath Palihawadana¹, Derek Bridge², Preeja
Pradeep², Belen Diaz Agudo³, and Marta Caro-Martínez³

¹ Robert Gordon University, Aberdeen, UK
{a.wijekoon1,n.wiratunga, d.corsar1,
i.nkisi-orji,k.martin3,c.palihawadana}@rgu.ac.uk,

² University College Cork, Ireland
{derek.bridge,preeja.pradeep}@insight.org

³ Universidad Complutense de Madrid, Spain
{belend, martcaro}@ucm.es

Abstract. Explainable AI (XAI) can greatly enhance user trust and satisfaction in AI-assisted decision-making processes. Numerous explanation techniques (explainers) exist in the literature, and recent findings suggest that addressing multiple user needs requires employing a combination of these explainers. We refer to such combinations as explanation strategies. This paper introduces iSee - Intelligent Sharing of Explanation Experience, an interactive platform that facilitates the reuse of explanation strategies and promotes best practices in XAI by employing the Case-based Reasoning (CBR) paradigm. iSee uses an ontology-guided approach to effectively capture explanation requirements, while a behaviour tree-driven conversational chatbot captures user experiences of interacting with the explanations and provides feedback. In a case study, we illustrate the iSee CBR system capabilities by formalising a real-world radiograph fracture detection system and demonstrating how each interactive tools facilitate the CBR processes.

Keywords: Interactive XAI, Ontology-based CBR, Conversational AI

1 Introduction

Explainable AI (XAI) is needed to guide users in understanding AI systems and their decisions. XAI systems must be able to address a range of user explanation needs (such as transparency, scrutability, and fairness) and must do so in a manner that is relevant to a range of stakeholders. Moreover, a successful adaptation of XAI should generate personalised explanations that better align with end-user mental models and cater to their specific needs. An interactive XAI

* This research is funded by the iSee project. iSee is an EU CHIST-ERA project which received funding for the UK from EPSRC under grant number EP/V061755/1, for Ireland from the Irish Research Council under grant number CHIST-ERA-2019-iSee and for Spain from the MCIN/AEI and European Union “NextGenerationEU/PRTR” under grant number PCI2020-120720-2.

system naturally creates a convenient feedback loop between the user and the XAI system, which is valuable for gathering user feedback to inform the system about their satisfaction regarding their needs. This feedback can refine the AI system and its explanation capabilities, improving its performance, reliability, and trustworthiness.

It is evident that developing meaningful XAI systems with positive user experiences is a multi-faceted endeavour. Consequently, it is essential for an AI system that is looking to adapt XAI practices to learn from past experiences of successful XAI adaptations. Case-based Reasoning (CBR) caters to the need to learn from past experiences. Accordingly, this paper presents the tools and processes that create a CBR recommender for reusing explanation experiences.

iSee is a consortium of researchers who proposed the use of the CBR paradigm to capture the knowledge and experience of successful adaptation of explainability within AI systems. iSee reuses these experiences with AI systems that are looking for the expertise to build explainability in their AI systems in line with regulations such as a right to obtain an explanation in the EU [5]. This paper presents the interactive tools in the iSee¹ platform that facilitate the explanation experience creation and reuse. The primary contributions lie in introducing three essential tools to enable CBR processes:

- requirements capture tool, to formalise explanation requirements modelled using the iSee Ontology;
- explanation strategy recommendation tool, to find similar past explanation experiences using case representation and retrieval; and
- feedback generation for revision and retention, by creating conversational explanation experiences modelled using a behaviour tree-driven dialogue model.

We demonstrate the effectiveness of the above-mentioned tools in the iSee system by presenting a case study that involves a radiograph fracture detection system. The outline of this paper is as follows. Section 2 provides the background and related work and Section 3 presents the overall CBR paradigm for interactive XAI. We describe the interactive components of the iSee platform in Sections 4, 5 and 6 respectively. Section 7 demonstrates the case study. Finally, Section 8 offers some conclusions and future directions.

2 Related Work

The CBR paradigm has played a key role in the development of methods and tools for reusing experiential knowledge. The flexibility of CBR lends well to capturing expert knowledge, modelling generalisable solutions and subsequent adaptation for bespoke scenarios. A key advantage of this is the ability to model solutions as plans; a sequence of steps to achieve a specific goal given a list of resources and constraints. Plans offer a rich representation whereby knowledge

¹ <https://isee4xai.com/>

of conditions of success and failure can be stored [10]. For example, CHEF [9] maintains knowledge of case outcomes to prevent the repetition of erroneous recipe adaptations. Business processes [27], production systems [19] and treatment strategies [17] are further examples of experience-driven domains that are satisfied by planning solutions. Similarly, we leverage CBR principles in reusing XAI experiences. We formalise and capture explanation experiences knowledge within the CBR cycle in the form of explanation requirements and strategies that satisfy them. It allows AI systems to reuse past experiences in validating the AI decisions, identifying potential issues, and improving trust [12].

2.1 Interaction Modelling of XAI

Conceptual models have attempted to capture multiple facets of XAI with the common theme that stakeholders have variable needs that are addressed by explainability techniques [2, 14, 16]. Two key components of such models focused on in this paper are the explanation techniques (i.e. strategies) and interactive interfaces to address stakeholder needs [25, 24]. The authors of [25] and [2] designed generalised strategies consisting of multiple explainers that addressed the needs of several application domains. In contrast, the authors of [24] derived XAI strategies specific to the healthcare domain from expert users. The process included studies to learn explanation needs, the results of which were used by the researchers to curate XAI strategies. There are both data-driven and expert knowledge-driven methods that exist to curate XAI strategies, but the challenge remains with their reusability across domains.

The usability of these strategies is linked to interfaces that interact with the users to extract knowledge, understand explanation needs, present explanations using different modalities, and generate feedback. Conversation is a medium for implementing interactive XAI, offering an alternative to graphical or text-based user interfaces [6, 18, 25]. Conventionally, conversational interactions are formalised as dialogue models using Argumentation Frameworks such as AAF [3] or ADF [20]. Alternatively, dialogue models are graphically represented using State Transition Models (STM) [11, 18] or Finite State Machines (FSM) [15].

The lack of shared conceptual modelling across XAI strategies and interactions discourages interoperability and reusability. This paper proposes the use of a unified conceptual modelling technique to model both XAI strategies and interactions using Behaviour Trees (BT). BTs are a less frequent choice for XAI dialogue modelling, although they have often been used to model robot interactions [4, 26, 13]. The design of an interactive model using BTs is either knowledge-driven by domain-experts [13, 4] or data-driven [26]. In this paper, we take a knowledge-driven approach using interaction requirements extracted from domain knowledge and previous work [11, 18].

3 CBR driven Explanation Experiences

In iSee, the Case-Based Reasoning (CBR) cycle [1], retrieves, reuses, revises, and retains explanation experiences as cases (Figure 1-left). We describe an




Fig. 1. iSee CBR System (left) and the high-level case structure (right)

explanation experience as a snapshot that captures the adaptation of XAI within an AI system. Accordingly, it is multi-dimensional, describing: the attributes of the AI system; user groups and their explanation needs; the explanation strategy; and user experience feedback. We formalise explanation experience cases using the iSee Ontology (iSeeOnto) as shown in Figure 1-right.

This paper focuses on three processes within the CBR cycle: 1) knowledge capture to form the case description; 2) case retrieval to recommend candidate explanation solutions from past experiences; and 3) conversational feedback gathering for collaborative case revision and retention. During the knowledge capture stage, we use iSeeOnto to understand and formalise the requirements for building an explanation experience. As a first point of interaction, we capture these requirements through a user interface from a design user of the AI system who has a working knowledge of the system’s development and stakeholders. These requirements form the query to our case base of past experiences, facilitating retrieval of the most suitable explanation strategies. A selected strategy is then provided to end-users for feedback in the collaborative revision stage. The following sections describe the tools and their underlying interaction models developed within the iSee platform to facilitate each process.

4 Explanation Experience Requirements Capture

To create meaningful user experiences, we structure interactions to manage the acquisition of explanation requirements. iSeeOnto bears the burden of information provisioning by dictating permissible features and values for each attribute such that the knowledge extracted from the design user conforms to a formal structure. These attributes describe the explanation needs of stakeholders associated with the AI system and provide information used to recommend an executable explanation strategy that best satisfies those needs.

An interactive interface is designed to capture knowledge from the design user and its driven by four ontologies: *AI Model*, *Explanation Criteria*, *User Group*,

and *Evaluation Metric*. High-level classes and relationships in these ontologies are depicted in Figure 2. Each class is further expressed by a taxonomy of classes or individuals. For instance, *AI Task* from the *AI Model* ontology is extended as a taxonomy with 50 hierarchical concepts, and *Intent* from the *User Group* ontology has 14 individuals identified from the literature. More detailed versions of each ontology can be found here¹.




Fig. 2. Ontologies associated with the explanation requirements capture (best viewed digitally in colour)

From a design user perspective, these ontologies formalise information to be displayed and acquired through the structured user interface called the iSee Dashboard. We propose that information requested by the dashboard should be provisioned by an individual familiar with an AI system and its stakeholders. We call this individual a design user and envision that they act on behalf of end-user stakeholders who will make routine use of, or have an interest in, the operations of the AI system. Inputs within the dashboard are divided into relevant sections to ease the design user’s cognitive burden and guide their provision of knowledge as explanation requirements. User input is also validated against the ontologies to ensure that only permissible values are captured, with support for users in the form of tooltips. As a result, once we capture these requirements from a design user, we build the description of a new explanation experience case. The design and implementation process of the dashboard is influenced by co-creation feedback from industry use case partners in the iSee project.

¹ <https://w3id.org/iSeeOnto/explanationexperience>

5 Explanation Strategy Recommendation

The next stage of the CBR process is case retrieval. Specifically, in iSee, this uses explanation requirements to find similar past cases that can recommend a candidate explanation strategy. Accordingly, this section describes the case representation, initial case base, and the interactive retrieval process.

5.1 Case Representation

Retrieval considers a subset of the knowledge acquired from the design user to form the query case. The attributes selected, along with an explanation strategy as the solution, forms the case representation presented in Table 1.

Table 1. Explanation Experience retrieval case representation and local similarities

Ontology	Case Attribute	Ontology Component	Similarity Metric	Solution
AI Model	AI Task	Class	Wu&Palmer [22]	-
	AI Method	Class	Wu&Palmer [22]	-
	Dataset Type	Individual	Exact Match	-
Explanation Criteria	Portability	Individual	Exact Match	-
	Scope	Individual	Exact Match	-
	Target	Individual	Exact Match	-
	Presentation	Class	Exact Match	-
	Concurrentness	Individual	Exact Match	-
User Group	Intent	Individual	Exact Match	-
	TechnicalFacilities	Individual Set	Query Intersection	-
	AIKnowledgeLevel	Individual	Exact Match	-
Behaviour Tree	DomainKnowledgeLevel	Individual	Exact Match	-
	User Questions	Individual Set	Query Intersection	-
Behaviour Tree	Explanation Strategy	N/A	N/A	✓

Case Description consists of 13 attributes. We select AI Task, AI Method, and Dataset Type classes from the AI Model ontology as attributes. AI Task and AI Method classes are expressed using their own taxonomies in iSee. The Dataset Type class consists of 5 individuals. Five Explanation Properties (Portability, Scope, Target, Presentation, and Concurrentness) were selected as case attributes. These case attribute values are inferred based on a set of rules (instead of asking the design-user) to ensure that the retrieved explanation strategy is compatible (for both implementation and execution) with the query case. AI and Domain Knowledge Levels and Intent attributes are selected from the

User Group ontology, where the attribute values are individuals of the respective classes. Finally, we consider Technical Facilities and User Questions classes where the case attribute value is a set of class individuals. For example, Technical Facilities are expressed using two sub-classes Audio Facilities and Visual Facilities and individuals such as *Speaker*, *Microphone*, *Touch Screen*, *Mouse*, and so on. Accordingly, a case can have multiple technical facilities; similarly, multiple user questions can express an explanation needs.

Case Solution is an explanation strategy composed of one or more explainers that address an explanation need. In iSee, an explanation strategy is modelled using Behaviour Trees (BT) and formalised using the Behaviour Tree ontology (high-level classes and relationships depicted in Figure 3).




Fig. 3. Behaviour Tree ontology in iSee

A BT is a conceptual model that formalises the behaviours of an actor in a given environment [7]. In addition to standard nodes and navigations (detailed in Section 6), we define and implement several specialised nodes to model explanation strategy behaviours. These include Composite Nodes *Variant*, *Supplement*, *Replacement* and *Complement* that model the relationships between multiple explainers or multiple presentations of an explanation. These are defined and formalised in the iSeeOnto. An example explanation strategy that satisfies two intents using three explainers is depicted in Figure 4. This explanation strategy can be interpreted as follows:

- If the user indicates *Transparency* as their intent,
 1. execute the Integrated Gradients explainer and show the explanation;
 2. afterwards, if the user indicates that they would like to *verify* (i.e. Variant Node) the explanation using a different explainer, execute the Nearest Neighbour explainer and show the explanation.
- If the user indicates *Performance* is their intent, execute the AI Model performance explainer and show the explanation.




Fig. 4. Example explanation strategy modelled using Behaviour Trees

5.2 Case Base

The iSee case-base currently consists of 12 *seed-cases*, captured from the literature; each describes the experience of adapting XAI within AI systems with user evaluation. This is a filtered list of cases from a literature review of 50 peer-reviewed papers to include only those who proposed reusable explanation strategies. All seed case explanation strategies consist of a single explainer addressing an explanation need (i.e. intent) of one or more stakeholders. We continue to add seed cases from the literature. In addition, we expect with time, the case base will grow by retaining new explanation experiences with more complex explanation strategies created within the iSee CBR platform.

5.3 Case Retrieval

Given a case with query attributes populated by iSee ontology classes or individuals, the retrieval task is to find explanation strategies from its nearest neighbours. We assign a local similarity metric for each attribute, as shown in Table 1. The details of those local similarity metrics are as follows.

Wu & Palmer (WP) is a taxonomy path-based similarity metric originally implemented for calculating word similarities. For AI Task and AI Method case attributes, we use the CloodCBR implementation [22] where, given a taxonomy, it calculates the similarity between two classes by considering the depths of each class from their least common subsumer.

Query Intersection (QI) is applicable for attributes where the data type is a set of ontology individuals like in Technical Facilities and User questions. Given a set of individuals from the query, s^q , and a case, s^c , it calculates the similarity as the intersection between two sets normalised by the length of the query set as $\frac{|s^q \cap s^c|}{|s^q|}$ where $|\cdot|$ indicates the size.

Exact Match (EM) similarity indicates a string match. This is applied both for case attributes that are ontology individuals and classes.

We formalise a case c as a list of N query attributes (a_i) and a solution (s) as in Equation 1. A query case q is a case where the solution s is empty ($s = \emptyset$).

$$c = [a_1, \dots, a_N, s] \quad (1)$$

$$\begin{aligned}
global_sim(q, c) &= \frac{1}{N} \sum_{i=1}^N local_sim(a_i^q, a_i^c) \\
local_sim &= \begin{cases} WP & \text{if } a_i \in [\text{AI Task, AI Method}] \\ QI & \text{if } a_i \in [\text{Technical Facilities, User Questions}] \\ EM & \text{otherwise} \end{cases} \quad (2)
\end{aligned}$$

The similarity between the query case q , and a case c from the case base is calculated as the aggregation of local similarities as in Equation 2. Note for iSee retrieval case structure, $N = 13$ as in Table 1. iSee case retrieval is implemented using the CloudCBR framework [22]. It is integrated with the iSee Dashboard, where the design user interacts with it by retrieving top k (configurable) cases, exploring the design of recommended explanation strategies, and making manual revisions to a selected explanation strategy. Aggregation of local similarities is currently unweighted, but CloudCBR allows for a weighted aggregation should that prove more useful, when the platform matures.

6 Conversational Feedback for Revision and Retention

As discussed in Section 5.3, we use Behaviour Trees (BT) to represent explanation strategies, these being the solution parts of our cases. But, additionally, we use BTs to model explanation experience interactions. We use them for this purpose because of their many desirable properties [8, 7] and also to give compatibility with the way we model explanation strategies. The tree structure is made of different types of nodes that implement behaviours and navigation. Each node has a state that indicates if the execution of the node was a success or failure. Composite nodes control navigation and the leaf nodes implement specific behaviour (Action Nodes). There are also decorator nodes and condition nodes to control access and repetition of a sub-tree. The types of nodes in the iSee dialogue model and their functionalities are briefly discussed as follows.

Sequence Node can have one or more child nodes and child nodes are executed from left to right until one fails.

Priority Node can have one or more child nodes and child nodes are executed from left to right until one succeeds.

Condition Node performs a Boolean check, often used as the first child node of a composite node with multiple child action nodes. The Boolean check helps to control the access to all its siblings to the right. For example, figures 5 and 6 show two scenarios where setting the *value = True* lets us control the access to the sibling nodes. In iSee conversations, this will help to avoid repetition and improve execution efficiency.

Explanation Strategy Node is a custom composite node introduced for iSee that can dynamically plug and play explanation strategies as the conversation progresses. It can be seen as a placeholder to be *replaced* when the specific explanation strategy is made available through the retrieval process (see Section 5.3).




Fig. 5. Condition node in a sequence sub-tree




Fig. 6. Condition node in a priority sub-tree

Evaluation Strategy Node is a custom composite node introduced for iSee that is a placeholder to *append* evaluation metrics (i.e. lists of questions) as the conversation progresses to multiple intents.

Action Node implements a specific behaviour. For example, in the iSee dialogue model, it will be behaviours of a chatbot in the format of *the chatbot prompting the user with an utterance, waiting for a response and analysing the response*. Based on the response, the business logic will determine its status as failure or success which helps the parent composite node to decide which node to navigate and execute next. iSee interactions are implemented in three custom Action Nodes: Question-Answer Node, AI-Model Node and Explainer Node. A Question-Answer Node will pose a question to the user and wait for a response which decides the node status. It is utilised to implement Start, Persona and Evaluation sub-trees. The AI-Model Node encapsulates the business logic related to the AI Model execution and is used in the Explanation Target sub-tree. Finally, the Explainer Node executes an explainer algorithm to generate explanations for the user and is utilised in the Explanation Strategy sub-tree.

6.1 iSee Dialogue Model




Fig. 7. Explanation Experience Dialogue Model

An abstract BT of the iSee dialogue model is presented in Figure 7. Each child is an abstraction of a sub-tree that handles a specific conversational behaviour. The most high-level navigation control is a sequence node, which means each child node should be executed successfully to complete an explanation experience. How each child (i.e. sub-tree) defines success is left to the business logic of the sub-tree. A simple execution of the conversation would be from left to right with the following steps: **a)** start the interaction by greeting the user and receiving consent to proceed; **b)** establish the persona, based on knowledge levels; **c)** establish the explanation target, i.e. the data instance and its AI system outcome; **d)** establish the user’s explanation need by asking questions, and present explanations to answer those questions by executing the suitable explainers of the explanation strategy; this is repeated until the user has no other questions or the XAI system is unable to answer any more questions; **e)** evaluate the experience using the evaluation questionnaire; and **f)** complete the explanation experience conversation. A fine-grained BT of the iSee dialogue model is included here¹ where each action node is expanded to its sub-tree.

At the end of a conversation, feedback for the Evaluation Metric of the case (questionnaire) is gathered and formalised as an individual of the User Evaluation Result (see Figure 8) to complete the case. Once there are multiple end-user experiences completed, we envision that those *User Evaluation Result* individuals will be analysed by the design user. If the feedback indicates failing to address explanation needs or disagreement with the explanations provided, the design user can iteratively revise the explanation strategy using the retrieval interaction. Otherwise, if the feedback indicates user satisfaction, the case is retained in the case base as a successful explanation experience for future reuse.




Fig. 8. User Evaluation Result ontology in iSee

7 Case study: Radiology Fracture Detection (RFD)

AI-assisted fracture detection through radiograph analysis accelerates diagnosis and treatment, which is particularly crucial in emergencies or high-volume cases [21]. However, achieving performance beyond established benchmarks requires Machine Learning algorithms, such as Convolutional Neural Networks

¹ <https://isee4xai.com/bt-2/>

(CNNs), which are black boxes whose outcomes are difficult to explain. Explanations help healthcare professionals understand the rationale behind the detection, offering insights to support their decision-making. In a recent survey of 411 UK radiographers, the most popular trust-building features of AI systems were *indication of overall performance* and *visual explanation* [23]. Based on such evidence and co-creation with industry partners we demonstrate how iSee utilises the interactive tools to create explanation experiences for the RFD system stakeholders. We capture the explanation requirements of stakeholders using the iSee Dashboard and utilise it to find suitable explanation strategies from past cases using retrieval tools. Finally, iSee dialogue model is instantiated to create interactive explanation experiences and collect feedback from stakeholders and complete the RFD case.

7.1 Explanation Experience Requirements Capture

Figure 9 presents part of the explanation experience requirements capture process with a design user of the RFD system. In AI model settings (left screenshot) the design user is using the iSee ontology concepts to describe their CNN model that performs Binary Classification which has been trained using an Image dataset. Also, they describe the performance of the AI model using two metrics F1-score and accuracy. RFD system has two main stakeholders who are interested in explanations, Clinicians and Managers. Clinicians can have explanation needs that are related to *transparency* or *performance* of the AI system. Managers mainly inquire about the *performance* of the AI system. The right screenshot demonstrates how the requirements of the *Clinician* user group are being captured in the iSee Dashboard.

Fig. 9. Explanation experience requirements captured in the iSee Dashboard with a design user of the RFD system (best viewed digitally)

7.2 Explanation Strategy Recommendation

Figure 10 demonstrates the case retrieval tool. The design user has retrieved top k (k=3 in the example) cases where each card component refers to a recommended explanation strategy. A strategy card provides additional information on how well the case description matches the query (global similarity) and what explainers are in the strategy. The design user can use this information to *select* an explanation strategy (the second strategy from Figure 4 is selected) and there are tools linked to *view* and *edit* the selected strategy.




Fig. 10. Case retrieval: case solutions from the top-3 neighbours for the *Transparency* intent of the *Clinician* user group (best viewed digitally)

7.3 Conversational Feedback for Revision and Retention

Figure 11 presents an instantiation of the iSee dialogue model by a clinician of the RFD system. First, the clinician is greeted and they select the user group they identify with. Next, they are presented with a test instance and its AI system prediction (Explanation Target sub-tree) which can lead to different explanation needs. The clinician selects a question that indicates transparency intent, accordingly, the explanation strategy selected by the design user for the transparency intent is executed. First, they are presented with an Integrated Gradients explanation, and upon requesting verification with a different explainer they receive a Nearest Neighbour explanation. At this point, the explanation strategy recommended for the transparency intent is completely executed. Now the clinician indicates they want to know about the AI performance (i.e. performance intent). Accordingly, the chatbot executes the explanation strategy recommended to answer performance-related questions. When the XAI system cannot answer any more questions (or if the clinician indicates they have no other questions), the BT exits the *Repeat Until Success* loop to the evaluation strategy sub-tree. It presents the clinician with evaluation metrics linked to both transparency and performance intents and collects their feedback. At this stage, the conversation is concluded and creates an explanation experience instance.

Good Morning 🌟 Anjana Wijekoon!
I am the iSee Chatbot for the Jiva Radiograph fracture detection system. Would you like to proceed?

Yes

I would like to know a more about your knowledge of the system, can you select a user group from below that best describes you?

Clinician
AI Knowledge Level: No knowledge
Domain Knowledge Level: Expert

Here is your test instance:

And here is the outcome from the AI system: Probability of fracture is 93.88%.

Okay

Please select a question below to indicate your explanation need

How accurate is the AI system?

Here is an explanation from AI Model Performance Technique

Assessment Metric	Value
F1-score (macro)	79.0%
Accuracy	77.8%

Explanation Description:
This explanation presents the performance metrics of the AI System.

Okay

Would you like to answer a questionnaire to evaluate your experience with iSee?

Yes

Please select a question below to indicate your explanation need

Which feature contributed to the current outcome?

Here is an explanation from Integrated Gradients Technique

Explanation Description:
Subplot with two columns. The first column shows the original image and its prediction. The second column shows the values of the attributions for the target class.

Okay

Would you like to verify using a different explainer?

Yes, I would like another explanation

This explanation of how the fracture detection system works is satisfying.

I agree strongly

This explanation of how the fracture detection system works has sufficient detail.

I'm neutral about it

This explanation of how the fracture detection system works is useful to my goals.

I agree somewhat

I am confident in the fracture detection system. I feel that it works well.

I agree strongly

I am wary of the fracture detection system.

I disagree somewhat

Thank you for using iSee! See you again soon!

Fig. 11. Explanation Experience of a clinician (best viewed digitally)

8 Conclusions

This paper presented the interactive tools that drive the iSee CBR cycle to reuse explanation experiences. The iSee ontology-driven Dashboard captures explanation requirements from the *design users* of an AI system, which is then used by the interactive case retrieval tool to find explanation strategies from past similar experiences. The iSee dialogue model facilitates *end users* of the AI system to create explanation experiences and provide feedback which then can be used in revision and retention. These interactive tools were co-created with industry partners and we demonstrated the tools on one such use case for detecting fractures in radiographs. In the next steps, user studies are planned to evaluate generalisability and user acceptance. This will also expand the range of explanation experiences to build a stronger case base.

References

1. Aamodt, A., Plaza, E.: Case-based reasoning: Foundational issues, methodological variations, and system approaches. *AI communications* 7(1), 39–59 (1994)
2. Arya, V., Bellamy, R.K., Chen, P.Y., Dhurandhar, A., Hind, M., Hoffman, S.C., Houde, S., Liao, Q.V., Luss, R., Mojsilović, A., et al.: One explanation does not fit all: A toolkit and taxonomy of ai explainability techniques. *arXiv preprint arXiv:1909.03012* (2019)
3. Baumann, R., Ulbricht, M.: Choices and their consequences—explaining acceptable sets in abstract argumentation frameworks. In: *KR*. pp. 110–119 (2021)
4. Bouchard, B., Gaboury, S., Bouchard, K., Francillette, Y.: Modeling human activities using behaviour trees in smart homes. In: *Proceedings of the 11th Pervasive Technologies Related to Assistive Environments Conference*. pp. 67–74 (2018)
5. Cath, C., Wachter, S., Mittelstadt, B., Taddeo, M., Floridi, L.: Artificial intelligence and the ‘good society’: the us, eu, and uk approach. *Science and engineering ethics* 24, 505–528 (2018)
6. Chromik, M., Butz, A.: Human-xai interaction: A review and design principles for explanation user interfaces. In: *Human-Computer Interaction—INTERACT 2021: 18th IFIP TC 13 International Conference, Bari, Italy, August 30–September 3, 2021, Proceedings, Part II* 18. pp. 619–640. Springer (2021)
7. Colledanchise, M., Ögren, P.: *Behavior trees in robotics and AI: An introduction*. CRC Press (2018)
8. Flórez-Puga, G., Gomez-Martin, M., Diaz-Agudo, B., Gonzalez-Calero, P.: Dynamic expansion of behaviour trees. In: *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*. vol. 4, pp. 36–41 (2008)
9. Hammond, K.J.: Chef: A model of case-based planning. In: *AAAI*. vol. 86, pp. 267–271 (1986)
10. Hammond, K.J.: Case-based planning: A framework for planning from experience. *Cognitive science* 14(3), 385–443 (1990)
11. Hernandez-Bocanegra, D.C., Ziegler, J.: Conversational review-based explanations for recommender systems: Exploring users’ query behavior. In: *CUI 2021-3rd Conference on Conversational User Interfaces*. pp. 1–11 (2021)
12. Hoffman, R.R., Mueller, S.T., Klein, G., Litman, J.: Metrics for explainable ai: Challenges and prospects. *arXiv preprint arXiv:1812.04608* (2018)
13. Iovino, M., Scukins, E., Styruud, J., Ögren, P., Smith, C.: A survey of behavior trees in robotics and ai. *Robotics and Autonomous Systems* 154, 104096 (2022)
14. Langer, M., Oster, D., Speith, T., Hermanns, H., Kästner, L., Schmidt, E., Sesing, A., Baum, K.: What do we want from explainable artificial intelligence (xai)?—a stakeholder perspective on xai and a conceptual model guiding interdisciplinary xai research. *Artificial Intelligence* 296, 103473 (2021)
15. Le, N.T., Wartschinski, L.: A cognitive assistant for improving human reasoning skills. *International Journal of Human-Computer Studies* 117, 45–54 (2018)
16. Liao, Q.V., Zhang, Y., Luss, R., Doshi-Velez, F., Dhurandhar, A.: Connecting algorithmic research and usage contexts: A perspective of contextualized evaluation for explainable ai. In: *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*. vol. 10, pp. 147–159 (2022)
17. Lopez, B., Plaza, E.: Case-based learning of strategic knowledge. In: *European Working Session on Learning*. pp. 398–411. Springer (1991)
18. Madumal, P., Miller, T., Sonenberg, L., Vetere, F.: A grounded interaction protocol for explainable artificial intelligence. In: *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*. pp. 1033–1041 (2019)

19. Malburg, L., Brand, F., Bergmann, R.: Adaptive management of cyber-physical workflows by means of case-based reasoning and automated planning. In: Enterprise Design, Operations, and Computing. EDOC 2022 Workshops: IDAMS, SoEA4EE, TEAR, EDOC Forum, Demonstrations Track and Doctoral Consortium, Bozen-Bolzano, Italy, October 4–7, 2022, Revised Selected Papers. pp. 79–95. Springer (2023)
20. McBurney, P., Parsons, S.: Games that agents play: A formal framework for dialogues between autonomous agents. *Journal of logic, language and information* 11(3), 315–334 (2002)
21. Moreno-Garcia, C., Dang, T., Martin, K., Patel, M., Thompson, A., Leishman, L., Wiratunga, N.: Assessing the clinicians’ pathway to embed artificial intelligence for assisted diagnostics of fracture detection. *CEUR Workshop Proceedings* (2020)
22. Nkisi-Orji, I., Palihawadana, C., Wiratunga, N., Corsar, D., Wijekoon, A.: Adapting semantic similarity methods for case-based reasoning in the cloud. In: Case-Based Reasoning Research and Development: 30th International Conference, ICCBR 2022, Nancy, France, September 12–15, 2022, Proceedings. pp. 125–139. Springer (2022)
23. Rainey, C., O’Regan, T., Matthew, J., Skelton, E., Woznitza, N., Chu, K.Y., Goodman, S., McConnell, J., Hughes, C., Bond, R., et al.: Uk reporting radiographers’ perceptions of ai in radiographic image interpretation—current perspectives and future developments. *Radiography* 28(4), 881–888 (2022)
24. Schoonderwoerd, T.A., Jorritsma, W., Neerincx, M.A., Van Den Bosch, K.: Human-centered xai: Developing design patterns for explanations of clinical decision support systems. *International Journal of Human-Computer Studies* 154, 102684 (2021)
25. Sokol, K., Flach, P.: One explanation does not fit all. *KI-Künstliche Intelligenz* 34(2), 235–250 (2020)
26. Suddrey, G., Talbot, B., Maire, F.: Learning and executing re-usable behaviour trees from natural language instruction. *IEEE Robotics and Automation Letters* 7(4), 10643–10650 (2022)
27. Weber, B., Wild, W., Breu, R.: Cbrflow: Enabling adaptive workflow management through conversational case-based reasoning. In: Advances in Case-Based Reasoning: 7th European Conference, ECCBR 2004, Madrid, Spain, August 30-September 2, 2004. Proceedings 7. pp. 434–448. Springer (2004)