

CARO-MARTÍNEZ, M., LIRET, A., DÍAZ-AGUDO, B., RECIO-GARCÍA, J.A., DARIAS, J., WIRATUNGA, N., WIJEKOON, A., MARTIN, K., NKISI-ORJI, I., CORSAR, D., PALIHAWADANA, C., PIRIE, C., BRIDGE, D., PRADEEP, P. and FLEISCH, B. 2024. Building personalised XAI experiences through iSee: a case-based reasoning-driven platform. In Longo, L., Liu, W. and Montavon, G. (eds.) *xAI-2024: LB/D/DC: joint proceedings of the xAI 2024 late-breaking work, demos and doctoral consortium, co-located with the 2nd World conference on eXplainable artificial intelligence (xAI 2024), 17-19 July 2024, Valletta, Malta*. Aachen: CEUR-WS [online], 3793, pages 313-320. Available from: https://ceur-ws.org/Vol-3793/paper_40.pdf

Building personalised XAI experiences through iSee: a case-based reasoning-driven platform.

CARO-MARTÍNEZ, M., LIRET, A., DÍAZ-AGUDO, B., RECIO-GARCÍA, J.A., DARIAS, J., WIRATUNGA, N., WIJEKOON, A., MARTIN, K., NKISI-ORJI, I., CORSAR, D., PALIHAWADANA, C., PIRIE, C., BRIDGE, D., PRADEEP, P. and FLEISCH, B.

2024

© 2024 for the individual papers by the papers' authors.

Building Personalised XAI Experiences Through iSee: a Case-Based Reasoning-Driven Platform

Marta Caro-Martínez^{1,*}, Anne Liret⁴, Belén Díaz-Agudo¹, Juan A. Recio-García¹, Jesús Darias¹, Nirmalie Wiratunga², Anjana Wijekoon², Kyle Martin², Ikechukwu Nkisi-Orji², David Corsar², Chamath Palihawadana², Craig Pirie², Derek Bridge³, Preeja Pradeep³ and Bruno Fleisch⁴

¹Department of Software Engineering and Artificial Intelligence, Universidad Complutense de Madrid, Spain

²School of Computing, Robert Gordon University, Aberdeen, Scotland

³School of Computer Science & IT, University College Cork

⁴British Telecommunications

Abstract

Nowadays, eXplainable Artificial Intelligence (XAI) is well-known as an important field in Computer Science due to the necessity of understanding the increasing complexity of Artificial Intelligence (AI) systems or algorithms. This is the reason why we can find a wide variety of explanation techniques (explainers) in the literature, on top of some XAI libraries. The challenge faced by XAI designers here is deciding what explainers are the most suitable for each scenario, taking into account the AI model, task to explain, user preferences, needs and knowledge, and overall, fitting into the explanation requirements. With the aim of addressing this problem, the iSee project was conceived to provide XAI design users with supporting tools to build their own explanation experiences. As a result, we have developed iSee, a Case-Based Reasoning-driven platform that allows users to create personalised explanation experiences. With the iSee platform, users add their explanation experience requirements, and get the most suitable XAI strategies to explain their own situation, taking advantage of XAI strategies previously used with success in similar context. The iSee platform is composed of different tools and modules: the ontology, the cockpit, the explainer library, the Explanation Experiences Editor (iSeeE3), the chatbot, and the analytics dashboard. This paper introduces these tools as a demo and tutorial for current and future users and for the XAI community.

Keywords

Case-Based Reasoning, Personalised Explanation Experiences, Explainer Library, Evaluation Cockpit, Explanation Experiences Editor, XAI Chatbot, XAI Ontology

1. Introduction

Nowadays, Artificial Intelligence systems (AI) help us overcome many daily tasks, in critical and challenging domains, such as healthcare, manufacturing industry or security. Since these tasks

Late-breaking work, Demos and Doctoral Consortium, colocated with The 2nd World Conference on eXplainable Artificial Intelligence: July 17–19, 2024, Valletta, Malta

*Corresponding author.

✉ martcaro@ucm.es (M. Caro-Martínez); anne.liret@bt.com (A. Liret); belend@ucm.es (B. Díaz-Agudo); jareciog@ucm.es (J. A. Recio-García); jdarias@ucm.es (J. Darias); n.wiratunga@rgu.ac.uk (N. Wiratunga); a.wijekoon1@rgu.ac.uk (A. Wijekoon); k.martin3@rgu.ac.uk (K. Martin); i.nkisi-orji@rgu.ac.uk (I. Nkisi-Orji); d.corsar1@rgu.ac.uk (D. Corsar); c.palihawadana@rgu.ac.uk (C. Palihawadana); c.pirie11@rgu.ac.uk (C. Pirie); d.bridge@cs.ucc.ie (D. Bridge); ppradeep@ucc.ie (P. Pradeep); bruno.fleisch@bt.com (B. Fleisch)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

have become crucial, the AI algorithms have been also developed to make them more accurate, which has led to the creation of more complex algorithms but, consequently algorithms that users find more difficult to understand [1]. To overcome this problem, eXplainable Artificial Intelligence (XAI) systems have been developed, trying to meet the necessity to increase user trust in AI and enhance their utility [2]. The research on XAI has grown by leaps and bounds due to high interest from the AI community and, as a result we can find a wide variety of explainability techniques (explainers) that can be applied to explain different AI models. The variety of explainers has also given rise to a new challenge: XAI design users find it difficult to know what type of explainers could be the best to implement, taking into account different factors, when they need to explain their own AI models and problems in specific contexts.

The iSee project tackles this problem, offering a platform that includes several tools for design users to find and build the best explanation strategies to apply in specific situations. The iSee platform is a Case-Based Reasoning (CBR)-driven platform, where users can share, retrieve, and reuse successful explanation experiences already applied in previous situations [3]. An *explanation experience* is understood as: (1) the set of needs, preferences, and constraints required to know when applying an XAI strategy; (2) the XAI strategy (the solution), which is a combination of explainers (answering users' questions) that we have represented in a modular fashion through Behaviour Trees (BTs); and (3) the evaluation of these strategies using feedback from end users.

In this work we present a tutorial/demo of the iSee platform. In the following sections, we first describe the fundamentals of the CBR cycle that drives the platform (Section 2). Later we guide the utilisation of the iSee platform through its components and tools (Section 3), and finally we draw some conclusions of the work (Section 4). At the end of this document (Appendix A), we provide readers with links to iSee online resources.

2. The iSee Platform CBR Cycle

A CBR system is composed of four main steps: retrieval, reuse, revise, and retain. Here, we briefly introduce the iSee CBR cycle process performed when a new use case is created in the platform. Design users will carry out the following process through the iSee platform: they will create their own use case where they are going to include all the information related to the specific scenario that they expect explanation for (AI model and AI task, user types, users' knowledge, users' intents, etc.). This is the **requirements capture step** that must be performed before the CBR cycle itself. After that, users will get a tentative list of recommended strategies, which are the XAI strategies successfully applied in the most similar previous situations (cases) that the CBR engine previously stored in its case base). This is the **retrieval step**. The requirement capture and the retrieval step are done in the *iSee cockpit*. Moreover, at this point, users have the option of creating a personalised strategy by clicking on a button in the cockpit. The button triggers the generation of an automatically adapted explanation strategy where explainers that address specific questions are collected from neighbouring cases based on the similarity of questions and we call this approach the **transformational reuse step** [4]. Thereafter, design users can choose one recommended strategy to apply to their use case (iSee suggests the best recommendation). However, they may need to edit this strategy if it is not fully applicable for

their specific use case (for example, the explainers in the strategy might not be applicable to the AI model that they want to explain). Even if the strategy is applicable, design users may want to change the recommended strategy according to their needs in a more personalised way. These tasks can be carried out through the *Explanation Experiences Editor (iSeeE3)* and constitute the **constructive reuse step**. Getting back to the cockpit, users can then design evaluation questionnaires which will be sent to end users after those end users query and test the explanation strategies. Once the use case requirements and the XAI strategy are ready, design users can publish it as a complete use case, and send the solution to end users for further use. This evaluation is done in the *iSee chatbot*. End users walk through the explanation results through a chat, also answering the evaluation questionnaire designed by the design user. Finally, design users can access the *analytics dashboard*, where they can view the results of the questionnaire. This is the **revise step** and may lead design users to change their strategy if the results are not successful, or may lead them to save the solution (**retain step**) in our case base for other design users to use it with their own XAI problems if the results are successful.

The iSee platform also includes other tools and components that help to carry out this process: the iSee ontology and the explainer library. The *iSee ontology (iSeeOnto)* defines the vocabulary describing an explanation experience. The cockpit and iSeeE3 define the requirements through the iSeeOnto definitions, and get the solutions in the retrieval and reuse steps using the semantic knowledge provided by iSeeOnto. The *explainer library*, which contains 70 explainers, provides users with the execution of the explainers, allowing them to view the explanation results during the reuse and the revise steps. The platform also allows users to include their own explainers using the cockpit through the description provided by iSeeOnto.

3. The iSee Tools

3.1. The iSee Ontology

The iSee Ontology (iSeeOnto)¹ [5] is the formalised representation of an explanation experience. It contains all the vocabulary and relationships necessary to define an explanation experience; we hope it may also help the XAI community guide the design of further explanation systems.

An explanation experience in iSeeOnto is described as a tuple $\langle D, S, O \rangle$. D defines the description of the situation that we need to explain. We included the AI model and AI task to explain, the end user profile, questions to answer, intentions, etc., and the explainability requirements. The explainability requirements define the desired type of explainers that needs to be included in the XAI strategy and which will frame the solution. The explainers are also defined in iSeeOnto through concepts well-known in the literature, such as portability, scope, concurrentness, data type, explainer implementation framework, among other features.

S is the explanation strategy that fulfils all the requirements included in the description D . An explanation strategy constitutes the execution of an explainer, or a set of explainers, that fit with the requirements in D . We have formalised an explanation strategy as a Behaviour Tree (BT) since it is a mathematical model that can execute explainers in a modular fashion.

¹Available at: <https://github.com/isee4xai/iSeeOnto>. Documentation available at <https://isee4xai.github.io/iSeeOnto/docs/explanationexperience-en.html>

Explainers and user questions will be in the leaves of the BT, while other composite nodes (internal nodes) will determine the execution of the explainers. iSeeOnto defines all of these terms and the relationships between them.

Finally, O is the outcome that we obtain when we evaluate S with real end users. Therefore, it represents the user satisfaction and ‘goodness’ with S for the problem described in D .

3.2. The explainer Library

The explainer library² is an API where we expose 70 explainers available “as a service”. This library is used in the iSee platform (specifically in the iSeeE3 tool and in the chatbot) every time we need to execute an explainer for design or end users to view the resultant explanation for their use case. The API may be used as an independent tool, since other XAI researchers or practitioners can make use of it, even without using the iSee platform.

iSee design users can even contribute to the explainer library: either because they want to include an explainer for other iSee users to use, or because they need that explainer as a part of the explanation strategy they want to shape in iSee. To do this, they need to include the explainer code in the iSee explainer library GitHub repository and later link all the explainer semantic information to iSeeOnto through the cockpit. This semantic knowledge, needed to perform the retrieval and reuse steps in iSee, is driven by the concepts modelled in iSeeOnto. In the cockpit, design users can see all the explainers already in the library (and all their semantic knowledge) and there is a button which the design users can select to fill a form where they are going to include the explainer knowledge.

3.3. The Evaluation Cockpit

The cockpit³ is the tool where design users can access the main functionalities of iSee. First users need to create an account on iSee, and login in. Once they have an account, they can start using the cockpit. In the cockpit, design users can contribute to the explainer library, as we have seen in Section 3.2, but the main task to perform in the cockpit will be the creation of an explanation experience scenario linked to a specific use case. To do that, users will perform the following process. They will create a use case: they will open a pop-up where they will indicate the name, the domain and the goal of the use case. Once done, a new screen appears, where the users fill a form (that is going to be introduced gradually). This form is the place where they include all the explanation requirements, which are divided into three sections: the AI model settings, the AI model upload, and the user personas. It is important to note that the form is guided by iSeeOnto, i.e., users will select an option in most of the fields regarding the concepts defined in the ontology. This is going to facilitate the retrieval step since iSee can see the similarities between the user requirements and previous cases.

The AI model setting section includes the AI task to explain, the AI method to explain, the data setting (dataset type, data type, number of features and number of instances), and model performance values (metric type and metric value). In the AI model upload section, users can include the model via an API URL, or via a model file. In this case, users need to include the file,

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

³Available at: <https://cockpit-dev.isee4xai.com/>. Code available at <https://github.com/isee4xai/iSeeCockpit>

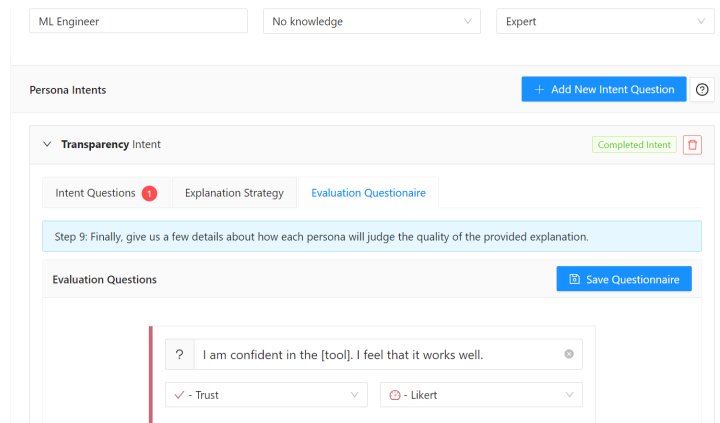


Figure 1: Cockpit screen to set the evaluation questionnaire for an intent and persona.

indicating its implementation framework, and a sample dataset file, where they will also indicate the data features, i.e., the type of features and possible values. Within the user persona section, the design users can create multiple personas. They have to indicate the persona name and their task domain and AI knowledge levels. Moreover, for each persona, they will add as many questions as they need from a list of questions related to an intent (user goal). These questions are the ones that user persona expect to get answered by reading the explanation generated by iSee. For each question, the design user will retrieve a recommended list of explanation strategies. To obtain the list of explanation strategies, we make use of CloodCBR⁴, a CBR tool that will compare the use case requirements provided by the design users –added through the cockpit– with the descriptions of previous cases. This list of explanation strategies will frame the XAI strategies of the most similar cases, from the iSee case base, to this specific use case. After retrieving the list, the design users will select one of them (after personalising it using the iSeeE3 tool should they wish to –see Section 3.4, or after personalising the strategy by clicking on a button to carry out the transformational reuse, where the explainers in the strategy are chosen be more suitable for the questions to be addressed). Then, they will include an evaluation questionnaire to evaluate that strategy for this specific intent and persona (see Figure 1). The questionnaire might be created by the design user or imported, using state-of-the-art questions that measure the end users’ satisfaction, ‘goodness’ or trust (for example, the Explanation Satisfaction Scale (Hoffman) [6]). The design users will need to complete this task for all their user personas, and their intents. Once everything is included, the design users will publish the use case. Doing this, they are allowing end users to evaluate the explanation strategy through the cockpit (see Section 3.5).

3.4. The Explanation Experiences Editor

The Explanation Experiences Editor (iSeeE3)⁵ [7] is the tool that allow design users to create their own explanation strategies, i.e., build their own BTs. It also allows them to execute the

⁴Available at <https://github.com/isee4xai/iSeeCloodCBR>

⁵Code available at: <https://github.com/isee4xai/ExplanationExperienceEditor>

resulting strategy on sample data instances: they can see the AI model prediction for their use case, and the explanations that are going to be retrieved by that explanation strategy. Although design users can create their own BTs from scratch using this tool, the iSeeE3 main functionality is to carry out the CBR constructive reuse step. The explanation strategies recommended by the cockpit might not be wholly applicable, or maybe the design users want to change some elements according to their preferences or needs. iSeeE3 include several functionalities to help them fix the strategies. Opening the iSeeE3 from the cockpit, the design users can access the recommended strategy (i.e. BT). They can modify it by hand, dragging components (composite nodes, or explainers). However, the reuse process might not be straightforward to design users. So there are two main functionalities within the tool to help users perform the reuse task more automatically: they can substitute individual explainers in the BT, or substitute whole subtrees. Both tasks can be done in a guided or automatic way.

The non-applicable explainer substitution consists of finding the most similar explainer (from our library) and replacing the explainer that is not applicable by the one that has been retrieved and that is actually applicable for the use case. In the guided mode, the user can pick the explainer substitution from a list of recommendations. In the automatic mode, users will click on a button and the non-applicable explainer will be replaced by the most similar and applicable one. The similarities between explainers and the applicability filtering are driven by the semantic knowledge from iSeeOnto. In the same way, users can substitute the whole tree. With the guided mode, users can pick one applicable BT (i.e. a BT where all explainers are applicable) from the list of the most similar BTs recommended by the tool. The recommendation list comes from the case base and it is calculated using a Levenshtein Edit Distance that also includes semantic knowledge from the explainers. With the automatic mode, the BT will be replaced by the one that is most similar to the one to replace, should this replacement be applicable. For both functionalities iSeeE3 also includes a form where users can indicate the type of explainers that they want in their BT. For instance, they might say that they want explainers that generate counterfactual explanations, or explainers that show explanations using heatmaps. Furthermore, the tool includes a button that will advise users whether the resulting BT structure is correct in terms of iSee rules.

3.5. The Chatbot and The Analytics Dashboard

Once a design user has built a specific use case (i.e. once the requirements and the explanation strategies are defined), end users will access the cockpit to evaluate its explanation strategy. The evaluation can be done through two tools: the chatbot and the analytics dashboard.

The chatbot performs the explanation strategies established by the design user. The chatbot is interactive: end users have to pick from possible answers to the questions that the chatbot asks them. First the end users will start by selecting the user persona that corresponds to them. Second, they will choose whether they want to upload a data instance or use inbuilt sampling method to select a data instance. Third, the chatbot will show the end user a prediction made by the AI model. The AI model is the one uploaded by the design user, and is executed through the dataset also provided by the design user. Then, the chatbot shows the questions established by the design user for that user persona (see an example in Figure 2). The end user will choose one question, and then the chatbot (following the execution workflow specified in the explanation




Figure 2: Explanation provided by iSee for a sensor anomaly detection use case.

strategy as a BT) will show the explanations one at a time. The user can choose to ask more questions (if any) or finish the process. In this case, the chatbot will allow users to answer the evaluation questionnaire that the design user has defined during the use case definition, finishing the process.

The answers from the end users for the evaluation questionnaire, combined with explanation strategy data, are available to the design users in the analytics dashboard. In this dashboard, they can see data such as the total number of interactions with the chatbot for this use case, and the number of interactions by persona. Also for each persona, the design users can view the number of interactions with each intent, the explainers performed, the answers to the evaluation questionnaire and the individual experience, i.e. the time spent by each end user in every step provided by the chatbot (explained in the previous paragraph). Together, interactions data allows design users to analyse how explanations were perceived by end users during their user experience.

4. Conclusions

The iSee platform is a CBR-driven tool that allow XAI design users to share, reuse and build personalised explanation experiences. Design users can: (1) determine the XAI problem that they need to solve for a specific use case; (2) obtain a personalised recommendation of previous successful explanation strategies for their problem; (3) build and personalise their own explanation strategies; (4) evaluate the resultant explanation strategy for their use case with real end users; and (5) share their explanation experience with other design users. In this work, we have described the iSee platform, and the tools and components that compose it. The main objective of the iSee platform is to be a useful tool for AI researchers and industry practitioners in order to encourage explainable and trustworthy AI. During the iSee project, we have completed three real world use cases successfully using the iSee tools: radiograph fracture detection, sensor anomaly detection, and telecom tasks blockers diagnosis. We also have studied about 10 use cases in other domains. Thanks to these evaluations, we have confirmed the iSee platform

utility for the XAI community. In the future work, we might look into ethical considerations to make sure users understand the risks involved when using AI systems or the iSee platform. Finally, we expect iSee to become more popular in the following months, so we can try our platform for different use cases, enriching the case base, and providing more different types of solutions, improving the platform itself as a consequence.

Acknowledgments

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 (with support from Science Foundation Ireland under Grant number 12/RC/2289-P2), for Spain from the MCIN/AEI and European Union “Next Generation EU/PRTR” under grant number PCI2020-120720-2, and for France from ANR under grant number 21-CHR4-0004-01.

References

- [1] A. Notovich, H. Chalutz-Ben Gal, I. Ben-Gal, Explainable artificial intelligence (xai): motivation, terminology, and taxonomy, in: *Machine Learning for Data Science Handbook: Data Mining and Knowledge Discovery Handbook*, Springer, 2023, pp. 971–985.
- [2] D. Gunning, D. Aha, Darpa’s explainable artificial intelligence (xai) program, *AI magazine* 40 (2019) 44–58.
- [3] A. Wijekoon, N. Wiratunga, K. Martin, D. Corsar, I. Nkisi-Orji, C. Palihawadana, D. Bridge, P. Pradeep, B. D. Agudo, M. Caro-Martínez, Cbr driven interactive explainable ai, in: *International Conference on Case-Based Reasoning*, Springer, 2023, pp. 169–184.
- [4] I. Nkisi-Orji, C. Palihawadana, N. Wiratunga, A. Wijekoon, D. Corsar, Failure-driven transformational case reuse of explanation strategies in cloodcbr, in: *International Conference on Case-Based Reasoning*, Springer, 2023, pp. 279–293.
- [5] M. Caro-Martínez, A. Wijekoon, J. A. Recio-García, D. Corsar, I. Nkisi-Orji, Conceptual modelling of explanation experiences through the iseonto ontology, in: *CEUR Workshop Proceedings*, volume 3389, CEUR Workshop Proceedings, 2023.
- [6] R. R. Hoffman, G. Klein, S. T. Mueller, Explaining explanation for “explainable AI”, in: *Proceedings of the human factors and ergonomics society annual meeting*, volume 62, SAGE Publications Sage CA: Los Angeles, CA, 2018, pp. 197–201.
- [7] M. Caro-Martínez, J. M. Darias, B. Diaz-Agudo, J. A. Recio-García, iSeeE3—The Explanation Experiences Editor, *SoftwareX* 21 (2023) 101311.

A. Online Resources

• [iSee webpage](#) • [iSee platform](#) • [GitHub of the iSee project](#) • [iSee overview video](#) • [iSee YouTube channel](#) • [Link to the chatbot to evaluate a sensor anomaly detection use case \(incognito mode required\)](#) • [Online guidelines to evaluation experiments](#)