

On the Explanation of Similarity for Developing and Deploying CBR Systems

Kerstin Bach

Department of Computer Science,
Norwegian University of Science and Technology
<https://www.ntnu.edu/idi>
kerstin.bach@ntnu.no

Paul Jarle Mork

Department of Public Health and Nursing,
Norwegian University of Science and Technology
<https://www.ntnu.edu/ism>
paul.mork@ntnu.no

Abstract

During the early stages of developing Case-Based Reasoning (CBR) systems the definition of similarity measures is challenging since this task requires transferring implicit knowledge of domain experts into knowledge representations. While an entire CBR system is very explanatory, the similarity measure determines the ranking but do not necessarily show which features contribute to high (or low) rankings. In this paper we present our work on opening the knowledge engineering process for similarity modelling. This work present is a result of an interdisciplinary research collaboration between AI and public health researchers developing e-Health applications. During this work explainability and transparency of the development process is crucial to allow in-depth quality assurance of the by the domain experts.

1 Introduction

Case-Based Reasoning (CBR) systems utilize previous experience in form of problem-solution pairs (cases) to solve new problems by matching the problem to its closest, most similar case (Aamodt and Plaza 1994). During the retrieval phase of a CBR system the case representation and similarity measures are crucial to find solutions that are most relevant to a given problem, while in reuse the modifications of the solution are undertaken to better suit the problem description.

While CBR is often described as an open box and explainable AI (XAI) method (Sørmo, Cassens, and Aamodt 2005; Johs, Lutts, and Weber 2018; Leake and Mcsherry 2005), the similarity assessment usually provides a shallow explanation in form of a similarity score. The interpretation of the similarity score, however, can be challenging and especially during the development of CBR systems discussions between experts and knowledge engineers need to be facilitated. From our experience, the more explanatory the reasoning process is, the better the knowledge representation and refinements from expert become. Especially considering the Knowledge Containers (Richter 1995; Ganesan and Chakraborti 2018), knowledge engineers have four different types of knowledge that interplay with each

other and hence influence the selection and ranking of cases. When the knowledge engineer or CBR expert does not have all available domain knowledge, collaboration with domain experts becomes increasingly important. Being able to explain the reasoning process within the domain with relevant cases helps to increase trust in the application by the expert as well as it leverages better evaluation in real-world setting and therewith detection of faults only.

In (Jaiswal and Bach 2019; Verma, Bach, and Mork 2018) we have introduced methods that support the knowledge modeling process for CBR systems. In this paper we focus on methods, especially visualizations, that explain the reasoning process – especially the similarity-based retrieval used in CBR systems. In the remainder of the paper we will present an overview of previous work of CBR and XAI. We will look into methods, applications and tools that contribute to enable explainable. In section 3 we will introduce a reference dataset for the remainder of the paper. We will present the current possibilities of explaining traditional machine learning methods before, in section 4 we show how similarity measures for the sample dataset can be defined. In section 5 we will introduce visualizations that allow a better understanding of the reasoning process. The final sections summarizes our work and gives an outlook on the next steps.

2 Related Work

The explanatory capabilities of CBR has been addressed by researchers throughout the life cycle of the field. Especially the work by Leake (Leake 2001) presents a general framework discussing issues that need to be addressed for explanations. Further on, (Sørmo, Cassens, and Aamodt 2005) present detailed explanation goals for a system: transparency, justification, relevance, conceptualization, and learning. Both works emphasize that not everything needs to be explained and that the context of an explanation needs to be taken into account.

Another aspect of explanations discussed in (Massie, Craw, and Wiratunga 2004) is that users on the one hand gain confidence in a system that provides correct results, but confidence is also improved when the decision making process is understood and deficiencies can be identified and resolved. This view on similarity measures and their role dur-

ing the retrieval will be addressed by our work later in this paper.

In previous work, the explanation of the case base content has been discussed by Smyth and McKenna (Smyth, Mullins, and McKenna 2000; McKenna and Smyth 2001) and they suggested visualisations of case base content and making changes to the case base explicit to the user. In more recent work, the authors of (Keane and Kenny 2019) present how CBR has been used to explain neural networks, which describes another branch of XAI.

While the theoretical concepts are important to move a field forward, their implementation in practice allows the research field to grow and attract others. CBR tools have been developed since the very beginning of the CBR research activities. The most general CBR tools developed and provided as bundled or open source software are COLIBRIS-tudio (and their predecessors COLIBRI, jCOLIBRI) (Díaz-Agudo et al. 2007), CBRworks (Schulz 1999) and its successor MYCBR (Stahl and Roth-Berghofer 2008). Furthermore there are more specific CBR tools targeting certain domains or case representations. For process-oriented CBR, the Collaborative Agent-based Knowledge Engine (CAKE) (Bergmann et al. 2014) has been introduced, while CREEK (Aamodt 2004) is a tool for knowledge-intense CBR and (B)EAR (Jalali and Leake 2015; 2016) focuses on the adaptation in CBR systems.

MYCBR¹, which is used and extended in this work, was developed by German Research Center for Artificial Intelligence (DFKI) and has been introduced as rapid prototyping tool for research and industrial applications. Recently the tool was generalized and provides a Rest API for more flexible interaction with the engine and its components (Bach, Mathisen, and Jaiswal 2016). Each component in MYCBR is explainable, which allows a deep integration of explanations in knowledge modelling, but also reasoning (Bahls and Roth-Berghofer 2007).

3 Example Dataset

CBR researchers have in the past very closely been collaborating with the health care domain due to its explanatory and transparent nature (Gonzalez, Lopez, and Blobel 2013). Since our work is also linked to patient-centered e-Health applications, we will introduce an open dataset from this domain to be used as a running example in this paper.

In the following we will describe different aspects of similarity modelling using an open dataset from the UCI Machine Learning Library (Dua and Graff 2017) containing 768 female patients of Pima Indian heritage. The dataset was introduced by (Smith et al. 1988) and has been provided by the National Institute of Diabetes and Digestive and Kidney Diseases to diagnostically predict whether or not a patient has diabetes, based on certain diagnostic measurements included in the dataset.

While the dataset describes the characteristics of patients included in the cohort, the main usage of the dataset is to classify whether a patient has diabetes or not. Figure 1 shows the value distribution within the diabetes dataset. Even when




Figure 1: Value Distribution in the Pima Diabetes Dataset




Figure 2: Benchmark of 8 classifiers on the diabetes dataset (kNN: k-Nearest Neighbour, SVC: Support Vector Machine, LR: Logistic Regression, DT: Decision Tree, GNB: GaussianNB, RF: Random Forest, GB: Gradient Boosting, XGB: XGBoost)

grouping them by outcome, the individual distributions are very similar, which shows that there is no apparent feature indicating the outcome.

As a reference point, we implemented standard machine learning approaches to carry out this task² and as it can be seen in Figure 2 with only basic tuning of the parameters the average accuracy only reaches about 0.78% for the best performing classifier on a 10-fold cross-validation.

Many of these methods would also allow to provide information about the feature importance for the classification, but this information still lacks detail on how the model is build and why certain features are chosen.

4 Similarity Modeling and Retrieval

The similarity modelling in a CBR system can vary from as simple as a kNN to knowledge-intense graph-based representations as presented in (Aamodt 2004). CBR tools, such as MYCBR that has been used in this work provide predefined similarity measure that cover mostly symbolic and numeric value ranges. In the example set of this paper we only have numerical values. The main similarity measure method of MYCBR is defining similarity according the local-global-principle (Richter 1995). Thereby similarity measures are defined as an amalgamation function, such as a weighted sum, which defines as local similarity measures the relationship of attribute values and the global similarity the weighted sum of the local similarities.

Defining Similarity Measures

Modelling similarity measure can be done automatically using neural networks (Gabel and Godehardt 2015), feedback

²The authors will make the code for training these classifiers available for the final version.

¹<http://mycbr-project.org>




Figure 3: Example of a polynomial similarity measure for the attribute *Age*: data distribution on the left and polynomial function on the right




Figure 4: Retrieval results from the diabetes CBR system

from users (Stahl 2005) or from the training data (Mathisen et al. 2019).

Figure 3 shows a data-driven or manual (expert-based) definition of a local similarity measure. The knowledge engineer can use data distributions (left of Figure 3) to define the characteristics of the similarity measure.

In the following step the global similarity measures – the weights for each particular attribute – are defined. This can either be driven by knowledge, derived from data or learned over time. Especially the reduction or expansion of the variables available in the dataset is part of this feature engineering process.

Retrieval

Once all similarity measures are defined, the case retrieval can be tested. The result of the CBR retrieval is usually a list of case-similarity pairs passed on to the adaptation engine and eventually presented to the user. For an example query to the diabetes dataset the result looks as shown in Figure 4.

There are different approaches on how many cases are selected for further processing. During development of CBR systems the knowledge engineer(s) often need to inspect the details of the similarity functions comparing the cases to verify correct behaviour.

5 Similarity Visualizations

In this section we present how retrieval insights can be presented to a knowledge engineer to gain understanding whether the similarity-based comparison is carried out as expected.

Figure 5 is an example of how the similarity scores leading up to the overall similarity presented in Figure 4. Each row of charts represents a comparison of a case from the case base to a query. The first row is the most similar case (highest ranked), followed by the second and third. The three charts are build the same way with the y-axis showing the

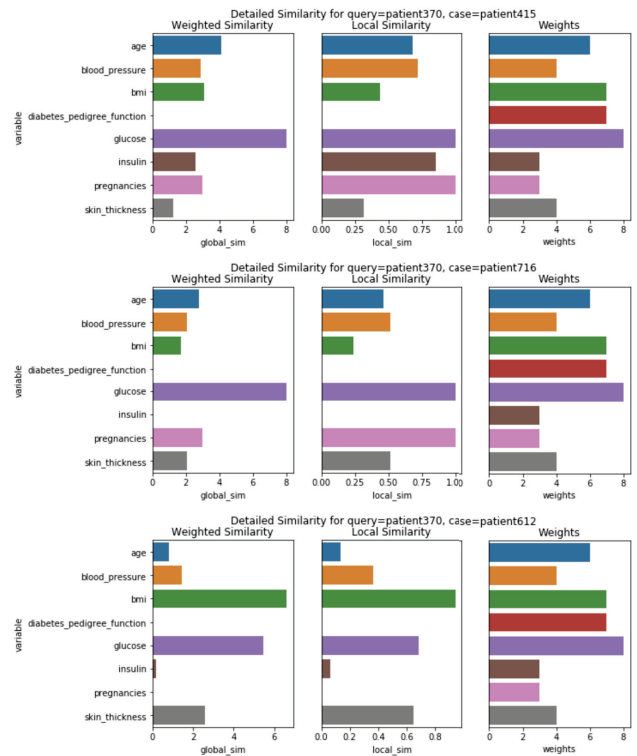


Figure 5: Visualization of weighted similarity, local similarity and weights

attributes and the bars the similarity. The left row shows the weighted similarity score, the middle the similarity scores from each local similarity measure and the third chart shows the weights.

In the particular example given, one can see that the second and third case which have a similarity score of 0.62 and 0.60 respectively reach that score through different attributes. For the second case mostly the glucose levels are matching while for the third a lower glucose, but higher BMI. Such insights are certainly important for the development phase of the CBR system.

6 Conclusion and Outlook

In this paper we presented on similarity measures can be explained during the development process of CBR systems. As our work is mainly carried out in interdisciplinary teams, the current transparency was not explanatory enough. Therefore we found visualisations that help to understand the retrieval and can be tested using the CBR engine. All visualizations presented have been implemented using the MYCBR Rest API and python, matplotlib and seaborn for data handling and visualization.

The next steps are to expand the visualizations towards the adaptation and case based evolution to gain better understanding when and how a CBR system is learning. With growing case bases, visualizing footprint cases and their provenance in the case base will also be in our focus.

References

- Aamodt, A., and Plaza, E. 1994. Case-based reasoning: Foundational issues, methodological variations, and system approaches. *Artificial Intelligence Communications*.
- Aamodt, A. 2004. Knowledge-intensive case-based reasoning in CREEK. In *Advances in Case-Based Reasoning, 7th European Conference, ECCBR 2004*, 1–15.
- Bach, K.; Mathisen, B. M.; and Jaiswal, A. 2016. Demonstrating the mycbr rest api. In *Workshops Proceedings for the Twenty-seventh International Conference on Case-Based Reasoning*, volume 2567. CEUR-WS.
- Bahls, D., and Roth-Berghofer, T. 2007. Explanation support for the case-based reasoning tool mycbr. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence (IJCAI 2007)*, 1844–1845.
- Bergmann, R.; Gessinger, S.; Görg, S.; and Müller, G. 2014. The collaborative agile knowledge engine CAKE. In *Proceedings of the 18th International Conference on Supporting Group Work, GROUP '14*, 281–284. New York, NY, USA: ACM.
- Dua, D., and Graff, C. 2017. UCI machine learning repository.
- Díaz-Agudo, B.; González-Calero, P. A.; Recio-García, J. A.; and Sánchez-Ruiz-Granados, A. A. 2007. Building cbr systems with jcolibri. *Science of Computer Programming* 69(1):68 – 75. Special issue on Experimental Software and Toolkits.
- Gabel, T., and Godehardt, E. 2015. Top-down induction of similarity measures using similarity clouds. In Hüllermeier, E., and Minor, M., eds., *Case-Based Reasoning Research and Development*, 149–164. Cham: Springer.
- Ganesan, D., and Chakraborti, S. 2018. An empirical study of knowledge tradeoffs in case-based reasoning. In *Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI-18)*, 1817–1823.
- Gonzalez, C.; Lopez, D.; and Blobel, B. 2013. Case-based reasoning in intelligent health decision support systems. In *PHealth 2013: Proc of the 10th Intl Conf on Wearable Micro and Nano Technologies for Personalized Health*, volume 189, 44. IOS Press.
- Jaiswal, A., and Bach, K. 2019. A data-driven approach for determining weights in global similarity functions. In *ICCBR-2019*, 125–139. Springer.
- Jalali, V., and Leake, D. 2015. Cbr meets big data: A case study of large-scale adaptation rule generation. In Hüllermeier, E., and Minor, M., eds., *Case-Based Reasoning Research and Development*, 181–196. Cham: Springer International Publishing.
- Jalali, V., and Leake, D. 2016. Enhancing case-based regression with automatically-generated ensembles of adaptations. *J. Intell. Inf. Syst.* 46(2):237–258.
- Johs, A.; Lutts, M.; and Weber, R. 2018. Measuring explanation quality in xcbr. In Michael T. Cox, P. F., and Begum, S., eds., *International Conference on Case-Based Reasoning*, 75 – 90. Springer.
- Keane, M. T., and Kenny, E. M. 2019. How case based reasoning explained neural networks: An xai survey of post-hoc explanation-by-example in ann-cbr twins. In Bach, K., and Marling, C., eds., *Case-Based Reasoning Research and Development*, 125–139. Cham: Springer.
- Leake, D., and Mcsherry, D. 2005. Introduction to the special issue on explanation in case-based reasoning. *Artif. Intell. Rev.* 24:103–108.
- Leake, D. 2001. Abduction, experience, and goals: A model of everyday abductive explanation. *J Exp Theor Artif Intell* 7.
- Massie, S.; Craw, S.; and Wiratunga, N. 2004. Visualisation of case-base reasoning for explanation.
- Mathisen, B. M.; Aamodt, A.; Bach, K.; and Langseth, H. 2019. Learning similarity measures from data. *Progress in Artificial Intelligence*.
- McKenna, E., and Smyth, B. 2001. An interactive visualisation tool for case-based reasoners. *Appl. Intell.* 14:95–114.
- Richter, M. M. 1995. The knowledge contained in similarity measures. In *Proceedings of the 1st International Conference on Case-Based Reasoning*, LNCS. Springer.
- Schulz, S. 1999. Cbr-works - a state-of-the-art shell for case-based application building. In *Proceedings of the 7th German Workshop on Case-Based Reasoning, GWCBR'99, Wrzburg*, 3–5. Springer-Verlag.
- Smith, J.; Everhart, J.; Dickson, W.; Knowler, W.; and Johannes, R. 1988. Using the adap learning algorithm to forecast the onset of diabetes mellitus. *Proceedings - Annual Symposium on Computer Applications in Medical Care* 10.
- Smyth, B.; Mullins, M.; and McKenna, E. 2000. Picture perfect: Visualisation techniques for case-based reasoning. In *Proceedings of the 14th European Conference on Artificial Intelligence (ECAI 2000)*, 65–72.
- Stahl, A., and Roth-Berghofer, T. R. 2008. Rapid prototyping of cbr applications with the open source tool mycbr. In *European conference on case-based reasoning*, 615–629. Springer.
- Stahl, A. 2005. Learning similarity measures: A formal view based on a generalized cbr model. In *Case-Based Reasoning, Research and Development, 6th International Conference, on Case-Based Reasoning, ICCBR 2005*, volume 3620, 507–521.
- Sørmo, F.; Cassens, J.; and Aamodt, A. 2005. Explanation in case-based reasoning—perspectives and goals. *Artif. Intell. Rev.* 24:109–143.
- Verma, D.; Bach, K.; and Mork, P. J. 2018. Modelling similarity for comparing physical activity profiles—a data-driven approach. In *International Conference on Case-Based Reasoning*, 415–430. Springer.