The Mining and Analysis Continuum of Explaining Uncovered

Martin Atzmueller and Thomas Roth-Berghofer

Abstract The result of data mining is a set of patterns or models. When presenting these, all or part of the result needs to be explained to the user in order to be understandable and for increasing the user acceptance of the patterns. In doing that, a variety of dimensions for explaining needs to be considered, e.g., from concrete to more abstract explanations. This paper discusses a continuum of explaining for data mining and analysis: It describes how data mining results can be analysed on continuous dimensions and levels.

1 Introduction

According to the CRISP-DM model [4] the data mining process consists of six phases: *Business Understanding* and *Data Understanding*, *Data Preparation*, *Modelling*, *Evaluation* and *Deployment*. These phases are ideally applied iteratively. In the evaluation phase the data mining models are checked and assessed by the user, before the models can be deployed: Often explanations for the complete models, or parts thereof are requested, e.g., for improving the acceptance of the patterns and their evaluation. Additionally, the mining process itself is a candidate for explanation, especially for inexperienced users. Appropriate explanation techniques in data mining and analysis are therefore crucial for an effective data mining approach; especially concerning *semantic data mining* and related approaches [2, 7], for which background knowledge provides further explanation capabilities.

Martin Atzmueller
University of Kassel, Knowledge and Data Engineering Group, Germany,
e-mail: `atzmueller@cs.uni-kassel.de`

Thomas Roth-Berghofer
German Research Center for Artificial Intelligence (DFKI) GmbH, Germany,
e-mail: `thomas.roth-berghofer@dfki.de`

This paper presents the *mining and analysis continuum of explaining (MACE)*; see [3] for a detailled discussion. The starting point of explanation is given by the final and intermediate results of the data mining step. Also, the specification of the data mining task itself can often be iteratively refined guided by appropriate explanation of the results. This also provides for a consistent documentation of the process and design decisions involved, e.g., in the form of semantic analytical reports, cf. [2, 7]. The recipients of the explanation sessions are the data mining engineer and the end-user. For both appropriate explanations are provided depending on the user role: While the end-user is mainly concerned with the evaluation and deployment phases of the cycle, the data mining engineer is involved in the whole process.

The rest of the paper is structured as follows: Section 2 describes the basics of explanation-aware design and computing. Section 3 describes general explanation goals and kinds. After that, Section 4 outlines the MACE, including explanation-aware mining and analysis, and the continuum. Finally, Section 5 concludes the paper with a summary and discusses further interesting options for future research.

2 Explanation-Aware Software Design and Computing

Software systems need the ability to explain reasoning processes and their results as such abilities substantially affect usability and acceptance. Explanation-aware computing (ExaCt) is the vision of software systems being smart in interactions with their users. Explanation-aware Software Design (EASD) aims at guiding software designers and engineers to a purposeful explanation-aware software system by making their designers and engineers explanation-aware. The long-term goal is to provide the respective methods and tools for engineering and improving the explanation capabilities. Here we focus on bringing explanation-awareness to data mining.

Explanations are in some sense always answers to questions, may the questions be raised explicitly or not. They enhance the knowledge of communication partners in such a way that they understand each other better. Explanations support humans in their decision-making [11]. In a general explanation scenario we distinguish three main participants [10]: the *user* who is corresponding with the software system via its user interface, the *originator*, i.e., the problem solver or 'reasoning' component, which provides the functionality for the original task of the software, and the *explainer*. Originator and explainer need to be tightly coupled to help the explainer provide knowledge about the inner workings of the originator.

As introduced above, we distinguish certain user roles in the data mining context: the end-user and the data mining engineer. The end-user considers the process as the overall originator, i.e., the data mining system is the only originator. The data mining engineer also receives input from this originator, but we can also embed distinct originators into the individual steps of CRISP-DM. Then, each of those also contains an explanation component for the individual steps that can also contribute to the (global) originator for the end-user.

3 Goals and Kinds of Explanations

For application development, there are two immediately useful classifications of explanation: Goals and kinds. In designing a software system knowing about kinds of explanations helps with structuring available knowledge and deciding which knowledge further is required for exhibiting certain explanation capabilities. Spieker distinguishes several useful kinds of explanations for knowledge-based systems [13]. *Concept Explanations* answer such questions as 'What is X?' or 'What is the meaning of X?'. *Purpose explanations* describe the purpose of a fact or object. *Why Explanations* justify a fact or the occurrence of an event. *Action explanations* are a special case of why explanations. They explain or predict the behaviour of 'intelligent systems'. *How Explanations* are similar to action explanations. They describe the function of a device without an actual context.

Explanation goals help software designers focus on user needs and expectations towards explanations and help to understand what and when the system has to be able to explain (something). Sørmo et al. [12] suggest a set of explanation goals addressing transparency, justification, relevance, conceptualisation, and learning. In [9], Roth-Berghofer and Cassens outline the combination of both, goals and kinds of explanations, in the context of Case-Based Reasoning, and examine the contribution of the four CBR knowledge containers for modelling necessary knowledge. In the following we take up this idea and cast it on the field of data mining.

4 The Mining and Analysis Continuum of Explaining

The Mining and Analysis Continuum of Explaining (MACE) provides different perspectives on the same problem. It considers different goals and kinds of explaining, presentation modes, levels of detail of explanation, knowledge containers, and privacy. In the following, we first describe the data mining foundations of the MACE, before we discuss its explanation dimensions.

4.1 Explanation-Aware Mining and Analysis

We regard the data mining system as originator, and provide explanation capabilities for each of the phases of the CRISP-DM model. The involved mechanisms can be summarised as follows: The input of the system is given by a (descriptive) specification of the process, the (source) data, and optional background knowledge. The system output is given by a data mining model, e.g., a set of patterns. The output is then accompanied by a "description" of the elementary mining steps, i.e., traces and logs of the respective events and steps of the process. The output can then be explained in terms of input data, additional background knowledge and intermediate results (trace). Additionally, setting up the specification itself is often a difficult task, for which appropriate explanation features are crucial.

- In the *Business Understanding* phase (concept) explanation helps inexperienced users getting accustomed to the domain, by structuring the relations between the concepts, and explaining the concepts in terms of their properties. Especially ontological knowledge is thus helpful for explaining concepts and properties.
- In the *Data Understanding* phase, important data elements need to be selected. Then, missing or redundant attributes can be added or removed from the data set. This can be accomplished by a concept explanation step. Furthermore, known correlations/dependencies between concepts can then be uncovered.
- *Data Preparation* and *Modelling* are strongly connected: Both can benefit from concept and purpose explanations, for configuring/specifying the mining task, and preparing the data accordingly. Additionally, how explanations consider the mining process and can be used for justification and transparency of the process itself; they show how the results were actually derived.
- In the *Evaluation phase*, the discovered models/patterns need to be assessed by the user. Therefore, they need to be interpreted and explained in a structured way using the concepts and/or contained patterns. The discovered patterns, for example, can be matched to semantic relations or more complex relations between these. Additionally, such knowledge provides a potential (explaining) context for the discovered patterns. The results of the evaluation can then be utilised for *task refinement*, e.g., for adapting parameters and/or method settings.

4.2 Explanation Dimensions (Continuum)

As outlined above, we distinguish different dimensions of explanation (Figure 1). In the following, we discuss them briefly in the mining and analysis context.

The user and/or application *goals* relate mainly to the kind of explanation. During data mining, a data-driven approach starts with the (intermediate/final) results of the mining step. Then, explanation is provided by analysing the trace of the system. Transparency of the results can be significantly increased by using contextual, why, how, or purpose explanations.

The *presentation dimension* of explaining needs to be performed in an appropriate way, e.g., using textual information, aggregation such as tables or visualisations for more aggregation and abstraction. The design issues involved here are also strongly connected to the *detail dimension*, since the level of detail needs to be reflected by the presentation options and the presentation modes need to be compatible with the detail level. In the continuum, the presentation dimension provides seamless drill-down/roll-up capabilities similar to OLAP [6] techniques connected with the detail dimension.

The MACE makes use of different *knowledge containers*, cf., [5, 8] that include explicit knowledge for explaining. We distinguish the containers *ontological knowledge (vocabulary)*, *pattern knowledge*, *instance knowledge*, and *context knowledge*.

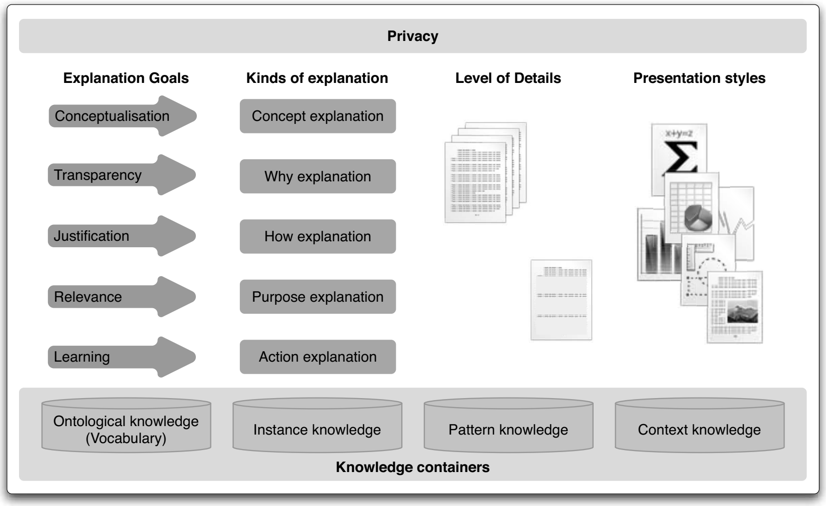

Fig. 1 Overview of the explanation dimensions

Whenever data is collected from heterogeneous sources, the aggregation of the data can reveal a lot more information than the single data sources. Privacy becomes an even more important issue with the availability and use of Linked (Open) Data.

In comparison to related application areas, e.g., case-based reasoning, the data mining and analysis domain provides for a more structured approach concerning the process, i.e., by applying the CRISP-DM cycle. In the individual steps of the process there are a variety of options for explanation, as discussed above. Additionally, the distinction between the 'inner' originators for the engineers and the 'outer' complete originator for the end-user, is also more present in the data mining context.

In practise, the proposed elements of the MACE need to be considered in a context of a specific data mining system. Additionally, the applied instantiation of the continuum also depends on the application domain. Both issues need to be considered when setting up the originator and explainer pair, and for arranging the match between them. Then, the utilisation of the instantiations of the dimensions depends significantly on the input context provided by the system, e.g., on the specification of the task, on the available trace information, and on the provided knowledge.

Since the knowledge containers are assigned both to the originator and the explainer, the specific knowledge containers can often be refined incrementally during the application of the system. While this is often easier considering the explainer, the extension and/or refinement of background knowledge applied by the data mining system is also possible. Several of the knowledge containers can often be reused 'as is' considering the originator, e.g., the ontological and instance knowledge containers. The pattern and context knowledge containers can usually be extended in the most flexible way, e.g., using Wiki-technology [1, 2].

5 Summary and Outlook

This paper presented a continuum of explaining for data mining and analysis: It described how data mining results can be analysed on several continuous dimensions and levels. We have described how the explanation options can be utilised in the standard *CRISP-DM* process model, and have briefly discussed the different goals and kinds of explanation in the context of the MACE.

For future work, we want to investigate ontological explanations in more detail, especially in the context of ubiquitous and social environments. Furthermore, appropriate tool support is necessary, especially regarding the presentation dimensions. Therefore, we want to investigate advanced explanation-aware presentation techniques in the context of the KNOWTA [1, 2] system, focusing on the concrete explanation-enhancing design issues.

References

1. Atzmueller, M., Haupt, F., Beer, S., Puppe, F.: Knowta: Wiki-Enabled Social Tagging for Collaborative Knowledge and Experience Management. In: Proc. Intl. Workshop on Design, Evaluation and Refinement of Intelligent Systems (DERIS), vol. CEUR-WS 545 (2009)
2. Atzmueller, M., Lemmerich, F., Reutelshoefer, J., Puppe, F.: Wiki-Enabled Semantic Data Mining - Task Design, Evaluation and Refinement. In: Proc. Intl. Workshop on Design, Evaluation and Refinement of Intelligent Systems (DERIS), vol. CEUR-WS 545 (2009)
3. Atzmüller, M., Roth-Berghofer, T.: Ready for the MACE? The Mining and Analysis Continuum of Explaining uncovered. Research Report RR-10-02, Deutsches Forschungszentrum für Künstliche Intelligenz (2010). ISSN 0946-008X
4. Chapman, P., Clinton, J., Kerber, R., Khabaza, T., Reinartz, T., Shearer, C., Wirth, R.: CRISP-DM 1.0: Step-by-Step Data Mining Guide. CRISP-DM consortium: NCR Systems Engineering Copenhagen (USA and Denmark) DaimlerChrysler AG (Germany), SPSS Inc. (USA) and OHRA Verzekeringen en Bank Groep B.V (The Netherlands) (2000)
5. Clancey, W.J.: The Epistemology of a Rule-Based Expert System: A Framework for Explanation. Artificial Intelligence **20**, 215–251 (1983)
6. Han, J., Kamber, M.: Data Mining: Concepts and Techniques, 2nd Edition. Morgan Kaufmann, San Francisco, USA (2006)
7. Kliegr, T., Ralbovsky, M., Svatek, V., Simunuk, M., Jurkovsky, V., Nerava, J., Zemanek, J.: Semantic Analytical Reports: A Framework for Post-processing Data Mining Results. In: ISMIS 2009: Foundations of Intelligent Systems, no. 5722 in LNAI, pp. 88–98. Berlin (2009)
8. Richter, M.M.: The knowledge contained in similarity measures. Invited Talk at the First International Conference on Case-Based Reasoning, ICCBR'95, Sesimbra, Portugal (1995)
9. Roth-Berghofer, T.R., Cassens, J.: Mapping goals and kinds of explanations to the knowledge containers of case-based reasoning systems. In: Case-Based Reasoning Research and Development, Proc. 6th Intl. Conf. on Case-Based Reasoning, no. 3620 in Lecture Notes in Artificial Intelligence LNAI, pp. 451–464. Springer Verlag, Heidelberg (2005)
10. Roth-Berghofer, T.R., Richter, M.M.: On explanation. Künstl. Intelligenz **22**(2), 5–7 (2008)
11. Schank, R.C.: Explanation Patterns: Understanding Mechanically and Creatively. Lawrence Erlbaum Associates, Hillsdale, NJ (1986)
12. Sørmo, F., Cassens, J., Aamodt, A.: Explanation in case-based reasoning – perspectives and goals. Artificial Intelligence Review **24**(2), 109–143 (2005)
13. Spieker, P.: Natürlichsprachliche Erklärungen in technischen Expertensystemen. Dissertation, University of Kaiserslautern (1991)