International Conference on Expert Systems for Development, Bangkok, Thailand, March 1994

CONTEXTUALIZED EXPLANATIONS

Brézillon P. J.

LAFORIA, Box 169, University Paris 6, 4 place Jussieu, 75252 Paris Cedex 05, France
Phone: +33 1 44 27 70 08 - Fax: +33 1 44 27 70 00 - E-mail: brezil@laforia.ibp.fr

Keywords: explanation process, context, knowledge-based systems, human-computer cooperation

Summary:
The aim of the paper is to fill the gap between theory and practice in the production of explanations by a system. One reason of this gap arises from the fact that a problem is often solved thanks to a cooperation between the user and the system, and both participants in the cooperation need explanations. Explanations essentially depend on the context in which the user and the system interact. Such contextualized explanations are the result of a process and constitute a medium of communication between the user and the system during the problem solving. We focus on the need to make the context notion explicit in the explanation process. We analyze explanation and context in term of chunks of knowledge. Then we point out what the contribution of the context to explanation is. An example, which is drawn from a real application, introduces what the problem is.

INTRODUCTION

A main line of research for introducing explanation capabilities in Knowledge-Based Systems (KBSs) supposes that explanations are used to transmit knowledge from the machine to a user and to improve learning ability of the latter. The lack of consideration for the end-user during the KBS design leads end-users to disregard these explanations. Conversely, there are KBSs that have been developed in an industrial context. These contain various types of explanation tailored for the needs of the user and the application, but do not explore all potentialities of explanations [BRE92]. There is a gap between theory and practice. One reason of this gap is that a problem generally is solved in a cooperative way by the user and the system. Then, an explanation is the result of successive transformations of a proposition that must be accepted by both the user and the system. Such explanations essentially depend on the context in which the user and the system interact.

Indeed, the context of the user-system cooperation takes into account different pieces of information like [MAS92]: The history of interaction with the user; The transaction history; The characteristics of the user; The user's intentions; The possible sources of ambiguity; The state of the system; The user's profile and system access allowed by the security.

Making explicit all these pieces of the context enables to: Tailor explanations to the users' needs; Simplify correctly a complex information for the user; Paraphrase; Structure the explanation; Manage counter-examples; Guide the research and the focus of attention; Resolve ambiguities; Correct misunderstanding; Help learning; Fill possible gaps; Reduce cognitive load; Develop a model of the "user-in-situation"; Make sure that the participants' mental models match; and Improve qualitative and quantitative performance of the explanation.

In this paper, we claim that a system must produce context-sensitive explanations to enhance its intervention in its cooperation with a user. A real-world example introduces the importance of the context notion in explanation in Section 1. We then expose what the role of explanation is according to the user's viewpoint in Section 2. We give a representation of the explanation process in Section 3. Section 4 exposes the notion of context that we retain. Section 5 discusses the relationships of context and explanation

1. A REAL-WORLD EXAMPLE

The example is drawn from a real application [BAU92]. First, we give information on the domain and the task of diagnosis in which the application has been developed. Then we give some of the various contexts and implicit knowledge that exist. In the following, we describe the two ways that have been chosen to implement the notion of context in an implicit way in our system. We end with a presentation of the advantages of making that notion of context explicit.

1.1 The domain and the diagnosis task

The French national power system is composed of controlled systems (Extra High Voltage substations) and control systems. Elements of the controlled system are lines and busbars that constitute a network. Elements of the control system are controllers. A controller links a busbar to a line or another busbar. It contains pieces of equipment like protective relays that act in parallel and a circuit breaker. Note that it is the appearance of the fault in the controlled system that triggers the protective relay, and its elimination that stops it. A protective relay detects the appearance of a fault in the controlled system and automatically sends a trip order to the circuit breaker after a delay. Actions of the equipment pieces are associated with signals that are sent to a central printer from which the operator diagnoses the events that occur.

The diagnosis concerns the equipment pieces in controllers that present a malfunctioning (e.g., the protective relay does not "see" the fault, or the circuit breaker does not open). The diagnosis is threefold:

(1) Determination of eventual malfunctioning of equipment pieces;

(2) Location of the fault in the controlled system;

(3) Detection of incoherence between fault characteristics and equipment functioning.

1.2 Contexts in the application

Context is not made explicit in this application. However, one can distinguish different contexts that involve: The diagnosis, the fault, the task of the controlled system (i.e., transport of electricity), the task of the control system, etc. Each of these contexts is composed of different contexts. For instance, diagnosis context depends on the contexts of: equipment, controllers, busbars (the global behavior of controllers that are connected to a busbar is judged in this context), and the substation.

1.3 Implicit knowledge

These contexts contain implicit knowledge as:
- Protective relays only operate when the fault exists,
- Protective relays "sees" effectively the fault,
- Transmission cables transmit correctly signals from equipment pieces to the central printer,
- The operations occurring in adjacent substations,
- The state of the overall network in which the controlled system is a node.

1.4 Context management

There are two versions of our expert system SEPT: an industrial one in Knowledge Craft, and a research one in a rule-based formalism. In the former version, contexts are represented by pre- and post-conditions, and screening clauses. In the latter version, contexts are represented by meta-rules that call rule packets in their action part when they are fired. We focus on the latter version hereafter.

A rule packet represents a set of conditions that hold in a given context. For instance, a rule packet describes the correct functioning of a circuit breaker. The system may check that rule packet for each controller in the substation. The meta-rule that calls the rule packet, permits to specify the context in which the system must check rules, i.e., specifies in which controller the circuit breaker is, and rules in the packet are fired for the circuit breaker of this unique controller. Thus, a context is given here by a set of instances of the variables.

We also use an interpretation notion with meta-rules. The notion of interpretation permits to select actions to execute in rules. It is thus possible to produce explanations at various levels of detail and abstraction. Now, the user manages by himself the interpretation under which he wishes to follow the system reasoning. Thus, the user acts on the context for obtaining the wished explanation.

1.5 Needs of context

Benefits for making explicit context in our application is of paramount importance:
- Make clear the knowledge that is represented (e.g., puting off screening clauses),
- Select relevant pieces of knowledge according to the current context of the diagnosis,
- Optimize diagnosis by making explicit the organization of the various contexts,
- Analyze of equipment behavior at a given level or at different levels of detail and abstraction,
- Ease the analysis of the signal flow and extract relevant information,
- Extend the use of the expert system: (1) at different substations; (2) for training purposes; (3) for its integration with other KBSs at Electricité de France.

We take in this application an approach that is opposed to the classical one, namely, the production of explanation from the system to the user. In our application, this is the user that built his explanation. The main reason is that the operator has a high level of expertise, even if this is a type of expertise that is different of those of the expert that has fed the expert system.

Indeed, there is a clear need of an intermediary solution where the explanation is produced as the result of a cooperation between the user and the system. Both of them have complementary competence.

2. THE EXPLANATION FOR A USER

Researchers agree about the goal of an explanation--a transfer of knowledge from the system to the user--since a long time. However, the finality of an explanation is frequently underestimated: An explanation is a process that allows the user to assimilate a new piece of knowledge. The keystone of assimilation is explaining the plausibility of new knowledge with respect to the knowledge that the user already possesses. The explanation process ends with the acceptance of the system answer by the user. This is a cooperative work where the answer given by the system is progressively adapted by both participants. Note that the user and the system must share goal(s) and knowledge.

An active participation of the system in the cooperation can only be achieved if the system is equipped with a number of capabilities for: Managing knowledge sources; Producing and tailoring answers to the user; and Acquiring information from the user and assimilating it in knowledge bases. Moreover, the user may need to take initiative in the explanation process by interrupting the system for: Pointing a part of the answer for which he has trouble; Verifying a particular element in the knowledge base; Asking further information; Correcting a false element in the system answer; and Providing spontaneous information to enable the system to evaluate the disagreement with him. Few systems now present such capabilities, although explanations constitute a medium of communication between the user and the system during the problem solving. This is the new challenge of the intelligent assistant system [BOY91a; BRE94].

3. THE EXPLANATION PROCESS

3.1 Chunks of knowledge

As noted by Draper [DRA88], the quality of an explanation essentially depends on the chunks of knowledge that are shared by the user and the system. Hereafter, we call CKs the chunks of knowledge that intervene during the explanation process. SK is the set of CKs that are under an explicit consideration by both participants during the explanation process at a given time.

The aim of an explanation is to make progressively explicit the links between the user's knowledge and a particular CK, called hereafter C_o, with which the user has trouble. The user will assimilate C_o if he integrates it in a SK in which all links between CKs will be made explicit by both participants. Thus, explaining is a process in which the system and the user cooperate.

3.2 An example

I may say to a person: "I heard a lion in my office this morning." It will be self-explanatory for a knowledgeable person. If the person is surprised (the person has some trouble with that C_o), I must develop my thought, saying: "I work in a university near a zoo that I can see from the window of my office. There are lions in that zoo. I often hear lions roar. It just was the case this morning." Here, C_o is introduced after various CKs are first presented.

The former person knows these CKs (that I work in a university near a zoo, that I can see it from my office and that there are lions in the zoo). However, I must make clear the CKs--the implicit knowledge--to share them with the person for explaining my first statement.

CKs are introduced in the SK progressively, beginning with CKs that may be accepted easily by the other. In the example, I begin with "I work in a University," and then "there is a zoo near my University."

3.3 Management of the CKs

The user may know or not C_o. Explaining a known C_o appears in two situations:

(1) The user does not know the link between C_o and the current SK,

(2) The user cannot choose between several links between C_o and the SK.

In these two cases, it is necessary to introduce new CKs in the SK.

For explaining a new C_o, the system first establishes which CKs of the current SK are relevant for C_o in the current context. Then, it determines which CKs are accepted by the user, makes clear to the user other CKs, and finally points out links between the CKs and C_o.

The system may introduce additional CKs to clarify some CKs in the SK. Then, it suggests links between shared CKs and C_o. The user may accept that construction or not, and eventually introduce other CKs to reinforce the SK according to his point of view. The explanation process ends when the new information is embedded in the shared SK.

Once a part of the SK only contains CKs that are shared (i.e., CKs are structured in a well-defined way), that SK may be compiled as a CK by both participants that will refer to the CKs afterwards as implicit and shared knowledge. Such a CK is similar to a pointer towards the subset of these CKs. Thus, the explanation process ends once the SK becomes a unique CK.

However, there is a difference when one of the participant is a computer system. For a human being, a SK only contains a compiled form of CKs, the steps of its building generally are lost. A computer system may have the SK, i.e., a compiled version of CKs, original CKs, and the steps that have been followed during interactions. Thus, a system may provide sometimes explanation more easily than a human being.

Another extension of this approach is a new insight of Grice's maxims. For instance, "Say something new" may be translated to "Introduce CKs that improve the SK building over CKs that do not." Such translations would permit Grice's maxims to be implemented easily.

4. A VIEW OF THE CONTEXT NOTION

4.1 Nature of the context

The SK under development constitutes a part of the context of the User-System Interaction (USI). The context also includes other types of knowledge: History, the system's state, a user model, the problem solving and the situation in which the problem is solved. Each type of knowledge defines a context by itself (e.g., the context of the user). Context now is one of the emerging challenges (e.g., see [BRE93]).

The context changes dynamically when the USI evolves, and CKs are added or removed of the SK during this evolution. Any participant (i.e., the user or the system) may introduce a CK in the USI context, first to establish a link with C_O, and second to share that CK with the other participant.

4.2 Role of the context

The context is a way to organize knowledge dynamically in view to facilitate changes in the SK at any time [BOY91b, DOL91]. Making explicit the USI context provides a mechanism that improves the efficiency of the cooperation between participants.

Several models have been proposed for representing a context like the traveling metaphor, the information theory and the conversation theory [MAS92]. For instance, a metaphor may provide information on a new context. It provides the context in which analogies and comparisons are made. It both limits and assists contextual understanding by sharing knowledge space between different situations and thus it may organize the whole knowledge base.

Making the USI context explicit also leads to a type of dynamic organization of the knowledge base. A SK is built for giving a unique meaning to a CK, even if the CK has several meanings. Thus, the meaning is tailored to the current situation in which the problem is solved and in a comprehensible way for the user because the user intervenes in the explanation process.

4.3 Importance of the context

Without context, the system may use CKs inappropriately. Capturing and using knowledge in context greatly simplified knowledge acquisition because knowledge provided by experts is always in a specific context and is essentially a justification of the expert's judgment in that context. Each acquired piece of knowledge is explained within the context of the other acquired pieces of knowledge and integrated with explicit explanations, i.e., with links with other CKs in the SK. Acquisition involves a CK and its links with the most relevant characteristics of the current situation in which its acquisition is needed, i.e., its context of use [WIC89]. This leads to a better organization of the acquired knowledge. Indeed, acquiring expertise involves developing a very systematic way of indexing information at the time of storage.

Specifying context, the system assimilates the acquired knowledge in a way that is well adapted to its knowledge and uses it directly. The process of assimilating new knowledge into a rich body of existing knowledge is a type of machine learning. Context-sensitive learning leads the system to organize knowledge bases according to clear criteria. Indeed, the system may acquire the way in which abstraction and generalization organize knowledge dynamically. There is not a unique knowledge organization. This is a weakness of these processes. Specifying context gives a solution to this problem.

4.4 Implementation of the context

Various formalisms of representation deal with the notion of context.

At a general level, a context is considered as a set of related propositions (descriptors) that describe a body of knowledge or a particular situation [SOW92] or through pragmatic rules. Management of contexts may use a mechanism of classification, a hierarchical organization [MAU92, DIE91] or around descriptors already used.

For instance, Maurer defines a context by the following statement [MAU92]:

Context
    name: LIGHT-BULB
    precondition:
        (SWITCH= CLOSED) & (LIGHT = OUT)
    correction: “Change the light-bulb”

In conceptual graphs, Sowa represents a context as a concept with one or more conceptual graphs nested inside the referent field. He now introduces an explicit representation of the context in his formalism [SOW92].

In rule-based formalisms, a context is a stable system of rules [MAS92]. It also corresponds to: An organization of rules; A special buffer-like data structure [BAR81]; A rule packet; Meta-rules that call rule packets [BRE92]. We also can find a notion of context at the rule level, mainly like screening clauses, and pre- and post-conditions.

In object-oriented formalisms, the context is a set of links between objects or attributes. At a lower level, the context is represented by a slot that links a concept with a state. This also may be a slot that despicts as a set of conditions upon functional units and links to environment in which this scheme can be identified.

In information systems, the notion of context plays a central role. A context may be a list of keywords, a relation between descriptors and referents [BOY91a], or a pathway in electronic documents [JAN91].

5. CONTEXT AND EXPLANATION

5.1 Influence of context on explanation process

A part of the context is the SK that is relevant for the goal of the USI. Thus, the USI context guides the explanation process.

The dynamic aspect of context implies that it is not possible to plan in advance the whole explanatory dialogue. That process reaches its goal when, after having exploiting all aspects of the shared context [CAR93], the user assimilates the SK including C_O with knowledge he already possesses. Such contextualized explanation account for a C_O in terms of factors that are not just causally relevant, but important to the explainee's future actions.

The context helps the user and the system to focus on the minimal number of relevant CKs at a time by making explicit the attention focus of the explanation.

More generally, the context provides a way to develop explanation strategies that are domain-specific, but are still quite domain-independent. This point is important because the expert knowledge is factually and tactically organized.

Supplying a context to a system gives senses to an important problem in explanation, namely the management of questions/answers. A proposed solution aims at organizing either questions or answers in a typology [GIL89]. Indeed, there is not a consensus on this topic. According to the context perspective, we note that context may enable to: Support the user in formulating and asking questions [GIB92]; Organize user's questions, making relationships among questions explicit [MOO89]; Determine the right meaning of a question in the current context [McK88]; and Facilitate the production of explanations of different types, at different levels of details and abstraction [BRE92]. The user may navigate easily in the explanatory hyperspace and thus avoid the "lost-in-space" problem because the context will restrict the search space.

We claim that accounting for context in explanation generation present a number of advantages such as [KAR94]: Give the right answer to the user's question; Relieve the user of having to formulate precise query; Provide concise and pertinent information immediately; Anticipate users' need for information.

5.2 Influence of explanation process on context

Conversely, the explanation process may be a mechanism for managing the USI context. The explanation process aims at the building of the SK in which an equivocal CK, C_O, may be assimilated by the user (i.e., the user may give a meaning to the C_O with respect to the SK). For instance, a spontaneous explanation may complete an information that is introduced in the SK for reinforcing the USI context [KAR92].

However, there is still much work to do in the study of the relationships between the explanation process and the context management

CONCLUSION

In this paper, we deal with the explanation process according to the point of view of the user-system interaction. This leads us to introduce a representation of that process, and discuss its relationships with context in this representation. Some implementation considerations of context are given, and an example drawn from a real application introduces the ideas that are presented in the paper.

We mainly discuss in this paper the role of the context in explanation. However, other tasks also are concerned with context: knowledge acquisition, automatic learning and knowledge assimilation, etc. The context also involves intelligent documentation, information retrieval, validation-verification, maintenance, all tasks that must be accomplished during the life cycle of a KBS. Some fields, like computational linguistics, already acknowledge the importance of the context notion [PAR91]. In our opinion, this is also a challenge for developing large reusable and shared knowledge bases. However, few systems deal now with such capabilities even though explanations may constitute a medium of communication between the user and the system during the problem solving.

However, there are a number of questions that have not yet received a clear answer. What are the elements of a context? We claim that context is knowledge. The SK in which a CK is embedded is knowledge. A CK finds a meaning in various SKs, i.e., various contexts, and may belong at different other SKs for explaining other CKs. One of our objectives is to clarify our position on this point. We also have not discus other aspects of context like time dependency. How could such an aspect be dealt with?

We think that the generation of contextualized explanations is a new challenge for the near future.

Acknowledgments
This work is partially supported by a contract LAFORIA-EDF-Airial and another contract between LAFORIA and CEA Marcoule.

REFERENCES

[BAR81] Barr A. and E.A. Feigenbaum (eds.), "The Handbook of Artificial Intelligence", William Kaufmann, Inc., 1981, Vol. 1, Chap III: Representation of Knowledge.

[BAU92] Bau D.Y. and P. Brézillon, "Model-based diagnosis of power-station control systems", IEEE Expert, February 1992, pp. 36-44.

[BOY91a] Boy G., "Intelligent Assistant System", Academic Press, London, 1991, Series Knowledge-Based Systems, N° 6.

[BOY91b] Boy G., "Indexing hypertext documents in Context", Hypertext'91, San Antonio, Texas, USA, December 1991.

[BRE92] Brézillon P., "Architectural and contextual factors in explanation construction", Proc. of the ECAI'92 Workshop on Improving the use of Knowledge-Based Systems with Explanations", Research Report 92/21, Laforia, University Paris 6, F-75252 Paris Cedex 05, France, June 1992.

[BRE93] Brézillon P., Proc. of the IJCAI-93 Workshop on "Using Knowledge in its context", Research Report 93/13, LAFORIA, Case 169, University Paris 6, F-75252 Paris Cedex 05, France, April 1993.

[BRE94] Brézillon P., "Design of an intelligent assistant system", Proc. of the International Conference on Expert Systems for Development, Bangkok, Thailand, March 1994 (this volume).

[CAR93] Carenini G. and J.D. Moore, "Generating explanations in context", International Workshop on Intelligent User Interfaces, Orlando, Florida, January 1993.

[DIE91] Dieng R. and A. Giboin, "Expert system explanations: the gap between research and industry", Proceedings of the IJCAI-91 Workshop on Explanation Generation for Knowledge-Based Systems, Sidney, Australia, August, 19911.

[DOL91] Dolmatova L.M., "Intellectual system of knowledge-based design for high pressure experimental mounting", Proc. of 1-st Moscow International HCI'91 Workshop, August 1991, pp. 249-259".

[DRA88] Draper S.W., "What's going in everyday explanation?", In: Analyzing Everyday Explanation, C. Antaki ed., Sage, 1988.

[GIB92] Giboin A. and C. Amergé, "Task/activity models as frameworks for identifying user's explanatory needs", Proc. of the ECAI-92 Wokshop on "Improving the Use of Knowledge-Based Systems with Explanations", Research Report 92/21, LAFORIA, June 1992, pp. 3-12.

[GIL89] Gilbert N., "Explanation and dialogue", The Knowledge Engineering Review, 1989, 3, pp. 235-247.

[JAN91] Jansen B., "Formal and narrative knowledge representation for explanations and justifications in knowledge-based systems", Technical Report TR-FD-91-02, CSIRO, Division of Information Technology, Sidney, Australia, May 1991.

[KAR92] Karsenty L. and P. Falzon, "Spontaneous explanations: a way to negociate concpts", Proc. of the ECAI-92 Wokshop on "Improving the Use of Knowledge-Based Systems with Explanations", Research Report 92/21, LAFORIA , June 1992, pp. 115-124.

[KAR94] Karsenty L. and P. Brézillon, "Cooperative problem solving and explanations", Journal of Expert Systems With Applications, 1994 (To appear).

[MAS92] Maskery H. and J. Meads, "Context: In the eyes of users and in computer systems", SIGCHI Bulletin, April 1992, 24(2): 12-21.

[MAU92] Maurer F., "Knowledge base maintenance and consistency checking in MOLKE/HyDi", EKAW'92, Lecture Notes in Artificial Intelligence, Springer Verlag, N°599, 1992, pp. 337-352.

[McK88] McKeown K.R. & R.A.Weida, "Highlighting user related advice", AAAI'88, Workshop on Explanation, 1988, pp. 38-43.

[MOO89] Moore J.D. and W.R. Swartout, "A reactive approach to explanation", Proc. of the 11th IJCAI, Detroit, MI, August 1989, Vol. 2, pp. 1504-1510.

[PAR91] Paris C.L., "The role of the user's domain knowledge in generation", The Computational Intelligence Journal, 1991.

[SOW92] Sowa J.F., "Representing and reasoning about contexts", AAAI'92 Spring Symposium on

Propositional Knowledge Representation, Stanford, CA, March 1992, pp. 133-142.
[WIC89] Wick M.R., "The 1988 AAAI Workshop on explanation", AI Magazine, Fall 1989, pp. 22-26.