

Using Knowledge in Its Context: Report on the IJCAI-93 Workshop

Patrick Brezillon and Suhayya Abu-Hakima

Advances in multidisciplinary research are leading toward a new generation of intelligent assistant systems (IASs) that are context sensitive (for example, see Boy [1991]). It is generally accepted that knowledge has a contextual component. However, even if its importance is acknowledged, this contextual component is rarely represented explicitly in available knowledge representation systems and is not used in subsequent processing of knowledge. Thus, there is a gap between what is known and what is done. Acquisition, representation, and exploitation of knowledge in context would have a major contribution in knowledge representation, knowledge acquisition, explanation, maintenance, documentation, learning, human-computer communication, and validation or verification. A computational capability to understand, represent, and reason about context will be valuable for, and of immense benefit to, many AI problems. For these reasons, a multidisciplinary group was invited to discuss the many issues surrounding context at the Workshop on Using Knowledge in Its Context, held prior to the Thirteenth International Joint Conference on Artificial Intelligence (IJCAI-93).

The workshop call for papers focused on computational principles and mechanisms for eliciting, modeling, and using the contextual aspect of knowledge as well as techniques and tools for building context-sensitive systems. The following series of questions were posed: What is con-

text? How can the notion of context be modeled? How can context be encapsulated in knowledge representation formalisms? What are the cognitive aspects linked to the notion of context? How does context influence

The Workshop on Using Knowledge in Its Context was held in Chambéry, France, on 28 August 1993, preceding the Thirteenth International Joint Conference on Artificial Intelligence (IJCAI-93). This article provides a summary of the discussions between the participants before (by e-mail) and during the one-day workshop. It is clear from these discussions that the notion of context is far from defined and is dependent in its interpretation on a cognitive science versus an engineering (or system building) point of view. In identifying the two points of view, this workshop permitted us to go one step further than previous workshops (notably Maskery and Meads [1992] and Maskery, Hopkins, and Dudley [1992]). Once a distinction is made on the viewpoint, one can achieve a surprising consensus on the aspects of context that the workshop addressed—mainly, the position, the elements, the representation, and the use of context. Despite this consensus on the aspects of context, agreement on the definition of context was not yet achieved.

human dialogue? How can the influence of context in human dialogue be taken into account in IASs? How can one exploit the notion of context in human-computer interaction (including assisting, explaining, or

clarifying interaction for a user)? What is the role of the notion of context in computer-supported cooperative work?

Early deadlines were stated to encourage participant interaction by e-mail before the workshop. The intent was to highlight the various positions of agreement and conflict that participants held on context. Thirteen position papers were selected and distributed to participants three months in advance of the workshop.¹

Preworkshop Topics

On the basis of the received papers, the program committee selected four topics on context for discussion before the workshop: (1) the position of context, (2) the elements of context, (3) the representation of context, and (4) the use of context. The *position of context* concerns what context is with respect to other entities, such as situation, behavior, point of view, meaning, relationships among agents, discourse, dialogue, and application dependency. Three examples illustrate the position of context: First, in explanation, the use of the context allows one to tailor a generated utterance to a user. Second, in knowledge acquisition, the explicit use of context limits the domain of validity of the acquired knowledge and indicates its correct moment of use. Third, in diagnosis, the explicit use of the context reduces the search to a smaller state space. A key question here is whether the process of achieving context—contextualization as introduced by William Edmondson and J. F. Meech (both of University of Birmingham, United Kingdom)—or the resultant state is of primary concern.

The elements of context include objects of the domain and the task at hand as well as the organization(s) of the elements, the mechanisms for managing them, and the manner in which these elements can manage domain knowledge. Further, what are the dimensions of the context, such as time and space (itself multidimensional), and the variables of the system (some, for example, could be

dependent on user profiles or models). Questions that arise from attempting to define the elements of context include the following: Is it sufficient to add a time dimension to account for context? Does context simplify or complicate the construction of a knowledge base? Is context nested in chunks of knowledge, or conversely, are chunks of knowledge nested in context(s)? Is context an object of the domain? Can we “decontextualize” knowledge (that is, move from one context to the next, as introduced by McCarthy [1993])? What are the links between the process of contextualization, as introduced by Edmondson and Meech, and decontextualization? What are the links between the external situation and the internal components of the context?

For the representation of context, one must examine the relation between the representation of context and the choice of formalism for representing and reasoning about a domain. Some believe that logic is the correct formalism, but others believe that semantic networks or conceptual graphs are better. Many formalisms are candidates. Some questions that arise from attempting to represent context include the following: What are the possible formalisms that seem to allow explicit representation of context? What are the comparison criteria of two formalisms with respect to context? What are the advantages and drawbacks of each of the possible formalisms? For a given formalism, which aspects of context are easily represented? Which ones are impossible to tackle? Such questions are important because it seems necessary to define what a context is with respect to the representation formalism that is used. However, it appears that the characteristics, the properties, and the like are dependent on the formalism that is chosen. Thus, it is difficult to have an abstracted view of what a context is, that is, at a conceptual level?).

A priori, there seem to be two uses for context: (1) to internally manage knowledge chunks and (2) to manage communication with a user. Some questions that arise in attempting to

typify the use of context are the following: Which representation formalisms allow reasoning in context? How do we use context to extract and present the relevant chunks of knowledge? Is context more important in some domains (for example, linguistics) than in others? Is the use of context definable as an activity or a process of contextualization? What are the links between internal and external management of knowledge?

Results of Participant Interaction

A key factor in studying any aspect of context is whether one’s view is that of a cognitive scientist or that of an engineer building a system to solve a particular problem. The *cognitive science view* is that context is used to model interactions and situations in a world of infinite breadth, and human behavior is key in extracting a model. The *engineering view* is that context is useful in representing and reasoning about a restricted state space within which a problem can be solved. On closer examination, one realizes that the engineering view is subsumed by the cognitive science view. Thus, participants from the two disciplines face similar problems in defining and using context and can share ideas in researching a solution.

Through participant interaction, it rapidly became clear that one cannot consider any of the four topics defined earlier. Moreover, it is difficult to speak of context independently of its use. (“We cannot speak about context out of context,” said Dan Suthers [University of Pittsburgh].) Also, participants are interested in context in different domains (logic, communication, knowledge acquisition, explanation, diagnosis, and so on.).

Thus, we reduced the objectives of the workshop to clarifying basic definitions and highlighting differences in the participants’ views on context. In the following subsections, we summarize the views and insights into context that were achieved by workshop participants.

Context and Its Use

Context is used in a number of domains. Indeed, it appears that there are as many contexts as uses of context because the notion of context seems inseparable from its use. Most of the papers in the proceedings (and elsewhere [Maskery and Meads 1992; Maskery, Hopkins, and Dudley 1992]) discuss aspects of context that are different. Thus, an abstraction of the various considerations of context at a conceptual level is difficult. One must first define what we want to do with a context in a given situation, even if we assume that we are always moving between contexts. (For example, must we speak of a discrete space of contexts or a continuous one?)

The use of context seems particularly important for a number of activities, including predicting context changes, explaining unanticipated events and helping to handle them, and helping to focus attention. Indeed, one can state that the activity in which one is engaged is itself the process of contextualizing data (Suthers). Context is what gives meaning to data, and “contextualization” is the process of interpreting data, transforming data into information (Edmondson and Meech). Moreover, the aspects of an agent’s context that are relevant to an agent’s activity can change as a function of the activity itself.

Cognitive processes are contextual. They depend on the environment (a part of the context) in which they are carried out. Thus, from a cognitive point of view, the context of communication in human-human dialogue can be defined as a set of transitory mental representations that are dynamically built by people during the interaction from the external situation that they perceive and from their knowledge, beliefs, and other characteristics (Béatrice Cahour and Laurent Karsenty, both of CNAM, France). These transitory representations are of the past dialogue, the task that they are performing, the spatiotemporal situation, and the psychosocial situation (representation of the other, of the social situation, and so on). They are selected according to their relevance in dia-

logue management, that is, to interpret what the other is saying and allow the other to interpret what we are saying (the process of contextualization, emphasize Cahour and Karsenty). Thus, the interlocutors assume that they share a large part of the dialogue context, but actually, they do not share exactly the same knowledge and characteristics and, consequently, the same activated context. Every participant then has his/her own vision of the assumed shared context, and these differences might not affect the communication process, although sometimes, the differences are too important, and a communication failure occurs. We can try to represent the context for man-machine communication needs, but we can never be sure of the context that a human interlocutor really has in mind during a dialog.

Situation and Current Context

The question of knowledge in context expands to its fullest scope in physically situated agents (Herbert Jaeger, University of Bielefeld, Germany). As a preliminary convention, one can propose to call *situation* the entire setting of a physical environment, the agent(s) therein, and the mental stuff within the agent(s). The *reasoning environment* is the mental stuff inside an agent, which at a given moment is the active background for its reasoning activity. The *situation image* is the portion of the reasoning environment that constitutes the agent's actual mental image of the situation. Note that the notion of a reasoning environment captures much of what many (but not all) might call context.

When trying to arrive at a comprehensive theory of situations, reasoning environments, and situation images, one faces some hard problems, three of which are the horizon, the system, and the typology problems described by Jaeger. The *horizon problem* states that situations and reasoning environments have no definite boundaries. The *system problem* states that the situation and the reasoning environment cannot be isolated but form an interdependent, incessantly interacting system. The

typology problem states that situations—and the reasoning environments along with them—are forever evolving in a qualitatively productive way and, therefore, cannot be reified and predescribed in a situation typology. Jaeger's approach to a representation format, *dynamic concept systems*, strives to satisfy the requirements that arise from these problems by integrating techniques of symbolic concept representation with insights from the system sciences.

Another important and related aspect of the context problem is establishing the current context. The current context is sometimes considered a combination of different contexts. For example, there is the context set of the application and that of the user. However, it appears that the most important context is that of the interaction, as shared by all participants in the interaction. All other contexts are either private contexts (Cahour and Karsenty) or parts of the interaction context.

Context and Reasoning

Context appears to be essential for all reasoning tasks and real-world problems. The use of the context in expressing reasoning can allow the exploitation of various forms of reasoning, such as nonmonotonic reasoning (within each context), reasoning about situations, approximate reasoning and a solution to the qualification problem, abstract reasoning, and metatheoretic reasoning (Fausto Giunchiglia, Istituto per la Ricerca Scientifica e Tecnologica, Italy). A context has also been considered a frame of reference for the interpretation of information both input to, and derived from, reasoning (Yves Lespérance, Simon Fraser University, Canada) (McCarthy 1993).

However, it is not clear if we must consider either reasoning within contexts (the reasoning environment) or the use of context for representing reasoning. Context can limit the domain in which the reasoning applies. For example, in reasoning about autonomous agents in cooperative, distributed problem solving, one must emphasize that contextual knowledge is crucial for solving real-

world problems (for example, controlling an autonomous underwater vehicle [Roy Turner, University of New Hampshire]). Here, the role of contextual knowledge is to help an agent behave quickly, automatically, and appropriately for its current problem-solving situation. Examples include (1) biasing situation assessment using top-down predictions about what the sensors are seeing, (2) modulating actions by setting behavioral parameters and automatically activating and deactivating goals based on the context, and (3) helping to choose actions appropriate for goals in the context.

Context can also be considered as a mechanism for managing, organizing, or reasoning about knowledge. One advantage of making context explicit in a representation is the capability of inference within and across contexts, thus explicitly making the change in reasoning across contexts (McCarthy 1993) (Giunchiglia). Context-sensitive reasoning can also be seen as a generalized form of case-based reasoning (Turner).

A central issue associated with the study of context is the relation between contexts and their organization. Contexts are related to each other: A context is defined in another context [McCarthy 1993]. Although contexts are related to each other in the basic definition, there are relations that can lead to an efficient management of movement from one context to the next. For example, using bridge rules, you can localize facts and reasoning to a single context before moving to another.

Context and Knowledge

Context provides a principled way to cluster, partition, and organize knowledge and its dimensions. This partitioning becomes relevant when contextual knowledge is integrated into problem solving, such as when navigating troubleshooting hierarchies (Suhayya Abu-Hakima, National Research Council of Canada).

One can reuse knowledge across contexts. This knowledge must be decontextualized, permitting one to abstract a piece of knowledge from contexts into a more general context

that covers the initial contexts.

However, it is not clear what the relationships between context and knowledge are at the level of their elements and at the level of their structures or organization. Lesperance argues that one needs to distinguish between having context-sensitive knowledge and knowing what the context is. For example, a robot might know that there is a pop can in front of it without knowing where it is, what time it is, and so on.

Components of Context

What is taken to constitute a context typically depends on the type of use you have in mind. In natural language understanding, it involves features of the discourse situation, such as speaker, hearer(s), time, and place as well as the dialogue history and mental state of the participants. In large knowledge-based systems, context is mainly seen as a mechanism for partitioning the knowledge base. In situated action, it covers actor, time, place, and objects involved; relevant background constraints on the situation (for example, structured lighting in robot vision); and background cultural assumptions, such as the individualization scheme (Lesperance).

The components of the current context that play a role in determining the context include much implicit, or background, knowledge. Making such implicit knowledge in the current interaction explicit permits you to include this contextual knowledge in bases of contextualized knowledge (B. Kang and Paul Comp-ton, both of University of New South Wales, Australia).

However, this task is not easy. The analogy of the *frame problem* was proposed (Vibhu Mittal and C. L. Paris, both of University of Southern California), which involves dynamically modeling side effects of the actions taken in the world by making corresponding modifications in the database that represents the state of the world as far as a system is concerned.

It appears that contextual knowledge changes in time by elaboration and shift. Thus, one must account for

both the *static aspect* (knowledge that remains constant throughout the interaction) and the *dynamic aspect* (knowledge that changes throughout the interaction) of context. The changing knowledge of a context and the movement between contexts would need to be managed by independent but related mechanisms.

Context and Representation

To select a formalism to represent context, one must also consider how to link contexts with knowledge and link contexts with other contexts. Various formalisms explicitly represent context in system building, including logic, rule sets, conceptual graphs and semantic networks (Abu-Hakima). Rule sets are easy to encode but hide some of the control knowledge that links contexts. Conceptual hierarchies, subsumed by conceptual graphs, are often used in diagnosis because they obviate control knowledge and localize contextual knowledge to a single concept linked to other concepts. Semantic networks are also subsumed by conceptual graphs and allow for a variety of relations to link concepts that are directly taken as contexts.

Other approaches have also been proposed for representing a context: a packet of knowledge pieces (Brezil-lon 1994) (somewhat related to conceptual graphs), a set of preferences (Suthers), a window (Abu-Hakima), an infinite and only partially known collection of assumptions (McCarthy 1993), a list of attributes, the product of an interpretation, and a collection of context schemas (Turner).

However, it is again not clear if the notion of context can be separated from the representation formalism. In addition, it is not clear if a single representation is adequate. One result of the workshop discussion is that formalisms for representing context are strongly dependent on modeling the requirements of context. Some of the key conflicts by participants at the workshop include some basic and divergent views: First, context can always be represented in well-defined domains. Second, context can partly be represented in narrow domains (for example, diagnosis). Third, context can never be represented from a

cognitive science view. Fourth, context can be represented without knowing what the context is. Fifth, context does not need to be represented because we only need to model the process of achieving context.

Some energetic discussion focused on the view that context does not need to be represented because we only need to model the process of achieving context, namely, contextualization (Edmondson and Meech).

Conclusions

Context is a notion that remains ill defined, with the workshop participants accepting no single definition. This message was clear throughout the workshop discussions and appears to be the case in writings on context in the published literature as well. One of the goals of the workshop was to clarify some basic definitions of context and point out alternatives. In this sense, this workshop brought new insights into the notion of context.

The discussions at the workshop helped to confirm that context can intervene for two types of use, (1) the management of knowledge and (2) the management of communication, which must be coupled to maintain consistency. Many alternatives and some divergent views were found among those attempting to model context in cognitive science and those attempting to use context in engineering systems. However, all participants agreed on the need for context in AI and its study in a multi-disciplinary group. For this reason, a follow-on workshop will be held at IJCAI-95.

Acknowledgments

The authors express thanks to the participants for their significant participation before and during the workshop. Thanks are also extended to all those who commented on a draft of this article, with special credit to Beatrice Cahour, Herbert Jaeger, Laurent Karsenty, and Roy Turner for contributing position paragraphs that were integrated into the text.

Notes

1. The proceedings are available from

Patrick Brezillon (Proceedings of the IJCAI-93 Workshop on Using Knowledge in Its Context, Research Report, 93/13, LAFORIA, University of Paris VI, Paris, France).

References

- Boy, G. 1991. *Intelligent Assistant System*. Knowledge-Based Systems Series, No. 6. London: Academic Press.
- Brezillon, P. 1994. Contextualized Explanations. In Proceedings of the International Conference on Expert Systems for Development, 228–233. Washington, D.C.: IEEE Computer Society.
- Maskery, H., and Meads, J. 1992. Context: In the Eyes of Users and in Computer Systems. *SIGCHI Bulletin* 24(2): 12–21.
- Maskery, H.; Hopkins, G.; and Dudley, T. 1992. Context: What Does It Mean to Application Design. *SIGCHI Bulletin* 24(2): 22–30.
- McCarthy, J. 1993. Notes on Formalizing Context. In Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence, 555–560. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.



Patrick Brezillon is a researcher at the National Center of Scientific Research in France. In 1983, he received his state doctorate in natural sciences at the Pierre and Marie Curie University.

The topic of his thesis concerned the mathematical modeling of the calcium metabolism as a self-oscillating nonlinear model. Now, he aims to merge the mathematical modeling with the AI approach and is concerned with topics such as cooperation, context, explanation, and incremental knowledge acquisition.



Suhayya Abu-Hakima has been a researcher in the Knowledge Systems Laboratory at the National Research Council of Canada since 1987. She previously worked for Bell Northern Research as a telecommu-

nications engineer. She is currently working on applications of AI to network management for seamless communication across heterogeneous networks. Her areas of research are diagnosis, explanation, multiagent systems, and context.