This paper appears in *Proceedings of the Second International and Interdisciplinary Conference on Modeling and Using Context (CONTEXT99)* (Trento, Italy, September, 1999) and in a forthcoming volume of the *Lecture Notes in Artificial Intelligence* series.

A Model of Explicit Context Representation and Use for Intelligent Agents*

Roy M. Turner
Department of Computer Science, University of Maine, Orono, ME 04469 USA
rmt@umcs.maine.edu

September 3, 1999

Abstract

Explicit representation of context and contextual knowledge is critical to intelligent agents. In this paper, we discuss our view of context and context-sensitive reasoning, based on several years of work on representing and using contextual knowledge. We describe our approach to context-sensitive reasoning, called *context-mediated behavior* (CMB), and discuss our experience related to reasoning in context in AI programs and our ongoing and future work in the area.

The context in which an intelligent agent operates profoundly affects how it behaves. This is both intuitive and has been supported by psychological and sociological studies (see discussion in [1]). This should be true for artificial agents as well. Indeed, it is difficult to imagine a definition of "appropriate behavior" that does not make reference to the context in which the behavior takes place.

Of course, AI programs have always taken their context into account to some extent. Usually, however, this has been done in an ad hoc way, without the designer paying explicit attention to context-sensitive reasoning, and without the programs having any explicit representation of their context or any clear sense of or access to their own contextual knowledge (cf. [2]). For example, AI planners (e.g., [3; 4]) create plans that must work in a particular task context, yet they do not represent the context as an object in its own right, nor do they explicitly identify the contextual aspects of planning knowledge. Observable features of the current situation constitute the program's view of its context, and its contextual knowledge is distributed in its operator preconditions and inference rules and, implicitly, in the assumptions encoded in the program itself.[1] Similar implicit context representation occurs in rule-based systems, neural networks, and other AI programs.

The result is that these AI programs are unable to capitalize on knowing what context they are in and how to behave in that context. They cannot do situation assessment, since they have no clear notion of what the space of possible situations might be, nor which classes of situations might have implications for how they are to behave. The programs are without the ability to truly reason about what contexts they are in and how they should behave while in them. Behavior is conditioned by

*This material is based upon work supported by the National Science Foundation under Grant Nos. BES–9696044 and IIS–9613646 and under contracts N0001–14–96–1–5009 and N0001–14–98–1–0648 from the U.S. Office of Naval Research. The content of the information does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

[1] Some planners do include meta-rules [5], but even here, only rudimentary attention is paid to representing the context as an object in its own right.

aspects of the current context, but not by the context as a whole. A program cannot conclude, based on its current knowledge of the world, "I am in context X", and then behave appropriately until the context changes. Instead, it must waste effort constantly deciding if its behavior is appropriate for the situation, for example by checking rule antecedents, goals/preconditions of operators, etc. Neither can a program easily learn important information about how to behave in a context, since it has no clear idea of what it means to *be* in that context. Further, it is difficult to acquire and maintain knowledge about how to behave in a context, since such knowledge will be distributed across many other pieces of knowledge (e.g., rules). This can lead to problems maintaining the consistency of the knowledge base.

For over ten years, we have investigated ways to explicitly represent contexts and contextual knowledge and to use those representations to control an agent's behavior. This work has taken place in the domains of medical diagnosis and intelligent agent (autonomous underwater vehicle) control. Currently, our research group is also examining context representation and use in multiagent systems and multi-modal interfaces. The approach developed, *context-mediated behavior* (CMB) [1], has strong ties to work in diagnosis, situation assessment, and case-based reasoning; indeed, we argue that identifying the current context *is* situation assessment, and that the process itself is a kind of diagnosis.

In this paper, we will first discuss context and our model of how context representations and contextual knowledge can be effectively used by intelligent agents. We next describe CMB, then discuss its development and use in two AI programs and our plans for its use in two other domains. Some conclusions and directions for future work are then presented.

1 Context Assessment

As has been pointed out repeatedly over the last several years at workshops, at conferences, and in on-line discussions focusing on context, the word "context" does not yet have a technical meaning that all researchers can agree with or even understand (e.g., [6]). Consequently, we need first to define what we mean by context in this paper, as well as provide some related definitions.

> **Definition 1:** In this paper, we use the term *world state* to mean the state of the agent's world at some particular time: i.e., all features of the world, including all objects in existence, their properties and internal states, and relationships between them.

> **Definition 2:** We define a *situation* to mean the portion of the world state that affects the agent in a significant way. The *perceived situation* is a set of features (which may or may not correspond to features in the real situation) that the agent believes belong to its current situation.

We realize that this appeals to the reader's intuition rather than being a formal or even a very exact definition. By this definition, the current situation includes all and only those features of the world that an omniscient viewer would realize affect the agent. So, for example, we would normally rule out the position of a satellite in its orbit as being part of the situation for, say, an autonomous underwater vehicle (AUV) on an oceanography mission, but would include that in the situation should the AUV's mission require it to communicate using that satellite.

The agent's knowledge about a particular situation is more than the sum of the features that are observable, even in principle, by an agent. An agent brings to a situation expectations about features it has not yet seen as well as predictions about future situations. It brings knowledge it has about the results of actions it took in previous, similar situations. It brings to bear knowledge of how it should act in the situation so that its behavior is appropriate. It has knowledge of how the situation is similar to and different from other situations.

An agent thus knows about *classes* of similar situations that share features and implications for how it should behave. We call such classes of situations *contexts*.

> **Definition 3:** A *context* is a distinguished (e.g., named) collection of possible world features that has predictive worth to the agent.

A context corresponds to a range of world states and situations. By "predictive worth", we mean that recognizing that the current situation is an instance of a known context is useful in helping the agent understand its situation and behave appropriately. Once an agent knows it is in a particular context—that is, that the current situation is an exemplar the context—then it immediately and relatively effortlessly knows a great deal about the situation, simply based on its *a priori* knowledge about the context. It can predict unseen features of the situation and how the situation may change over time. It knows what constitutes appropriate behavior in the situation.

Consider medicine. Doctors (and medical AI programs) do not as an end in itself perform diagnosis to identify what disease is present. There is no reason in principle that treatment could not be prescribed and predictions made about the patient's future health (prognosis) based solely on reasoning about the signs and symptoms, without recourse to explicitly identifying the disease(s). After all, the disease is not a real thing, in the sense of something present in the world that one can point to. It is just a name for a particular constellation of signs, symptoms, pathophysiological states, presence or absence of etiological agents (e.g., bacteria), and so forth. But the concept of "disease" provides the medical reasoner with several benefits. It clusters *predictive* knowledge about commonly-occurring and/or important recurring patterns in the world (i.e., the disease). Recognizing the presence of the disease allows the diagnostician to make rapid, accurate predications about the prognosis for the patient. More important, it allows appropriate treatments to be associated with the pattern; when the disease is detected, knowledge about the treatments immediately comes to mind. By explicitly recognizing the existence of these recurring patterns, the diagnostician can reason about them and compare them to one another. Thus the diagnostician can engage in differential diagnosis, using predictions from one disease to gather evidence to differentiate it from other candidate diseases [7]. Having explicit names for such patterns also allows doctors to communicate easily with others about the disease.

Diseases are just one kind of context.[2] There are a myriad of others, recurring patterns in the welter of the features of the world that an agent should recognize because they have predictive worth. For example, an autonomous underwater vehicle (AUV) should recognize such contexts as being in a harbor, operating under sea ice, and having low power. By understanding that these are important classes of situations, the agent gains the same advantages as does the diagnostician from knowledge about diseases. The context can serve to organize the agent's knowledge about predictions that can be made and behavior that is appropriate in the situation. Further, if the contexts are explicitly represented, they can allow the agent to do differential diagnosis to determine which among several possible contexts it is actually in. And explicitly representing contexts allows the agent to communicate with others about those contexts.

A major task facing an intelligent agent is *situation assessment* (e.g., [8]): given the features of the current situation it can see and its other beliefs about the situation, what context does this correspond to? We believe that situation assessment, which we could also think of as *context assessment*, is a process of diagnosis.

[2] More precisely, the presence of a particular disease in a patient is a kind of context.

In medical diagnosis, the history of the patient, signs and symptoms, and other features of the diagnostic situation all are used to diagnose the presence of a disease (or set of diseases). This is an abductive process that has been approached in a number of ways in AI (e.g., [9; 7]). One very effective approach was encoded in INTERNIST [7]. In this approach, findings *evoke*, or bring to mind, candidate disease hypotheses. For example, an abnormal chest x-ray might evoke lung cancer, tuberculosis, and possibly other diseases. That is not the end of the story, however; the diagnostician needs to assess how well each hypothesis fits the current situation. When two or more disease hypotheses conflict, the diagnostician also needs to determine which of them (if any) is indeed present in the patient. This can be done by *differential diagnosis*, the process of using predictions made by the hypotheses to differentiate between them.

We see this as a good model of context assessment, as well. Features of the perceived situation should evoke one or more contexts that are candidate assessments of the situation. Where the fit is sufficiently good, and where there are no competing contexts, then the process is complete. For example, when walking into a church on Sunday morning, with people sitting in the pews and a minister in the pulpit, the evoked context is "at a church service". Very few other contexts fit this situation. Immediately the person knows a wealth of things about the situation he or she has not observed, including predictions about the kinds of things the minister will or will not say, that there will be hymns sung, and so forth. He or she also immediately knows how to behave and how not to behave. Sitting down quietly and paying attention is appropriate; turning cartwheels down the aisles is not. This addresses Oztürk's [10] identification of two major functions of context, relevancy and efficiency. Contextual knowledge associated with a context is the knowledge that is relevant to the situations that context describes. The association of the knowledge with the contexts addresses efficiency by allowing knowledge appropriate for behaving in the context to be evoked along with the context.

Sometimes, however, there will be competing hypotheses about what the context is, just as there are often competing disease hypotheses in diagnosis. For example, suppose a person walks into a room at a university in which he or she often studies. There are a few people sitting around in a circle, with one person standing addressing them. Several contexts might be evoked in this situation, including: a class is going on; a meeting of a club is in progress; and a group of friends are talking. The person needs to differentiate between these competing contexts if he or she is to know whether it is appropriate to stay and study or to leave unobtrusively. Predictions based on knowledge about each context can help in this differential diagnostic process. For example, the person might look to see if the people sitting are taking notes, or if the person addressing them looks like a professor, etc. This all would happen quickly, but nonetheless, something very much like it would happen.

Contexts have the property of composability. Just as a patient can be suffering from several diseases simultaneously, the current situation can be characterized simultaneously by several contexts. In our example above, in addition to the contexts evoked by seeing the people listening to the speaker, being in the room itself forms a context, as does the fact that the person is (say) hungry, that he or she has work to do, etc. Each of these contexts or contextual aspects [11] carries with it predictions about the world (e.g., "this room is generally quiet") and suggestions for behavior (e.g., "find a place to work"). All of the contexts that fit the current situation can be composed, or merged, to give rise to the complete context assessment. For example, suppose an AUV is on a search mission, has low power, and there is a hurricane overhead. Knowledge about each of these contexts could be combined to yield a description of the current, perhaps novel, context that could then help the AUV behave appropriately.

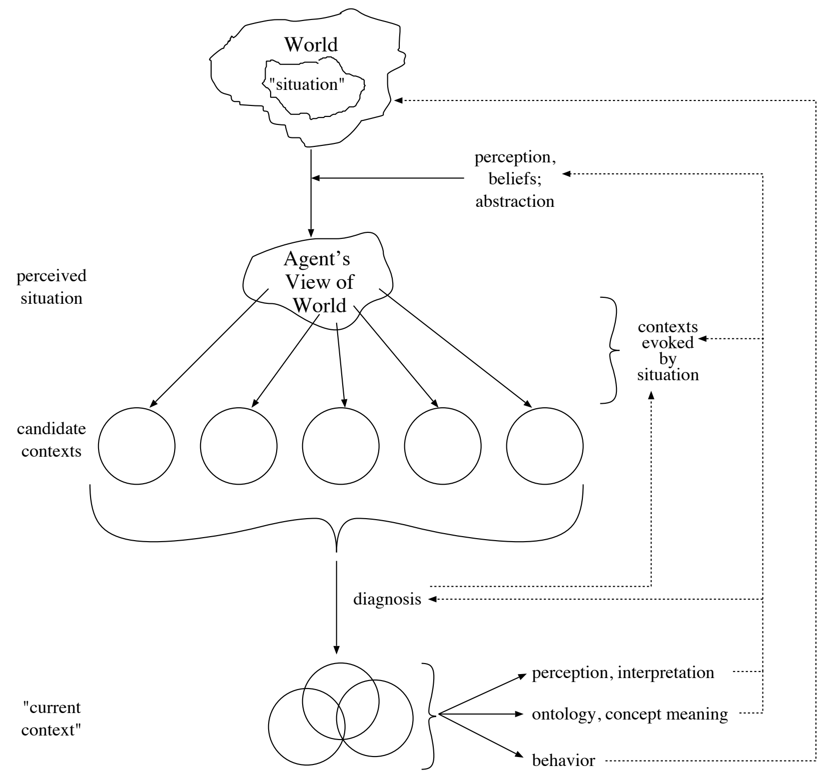

Figure 1: Overall process of determining the context.

Figure 1 summarizes our view of how an agent does context assessment and uses contextual knowledge. The agent's perceived situation is the product of its perceptions (including communication from other agents), its beliefs, and the process of abstraction and interpretation (e.g., recognizing that a computer is present from seeing its keyboard and display). Candidate contexts are evoked based on features of the perceived situation. A process of differential diagnosis then determines which of these contexts are the most appropriate characterizations of the current situation. The combination of these contexts gives a complete description of the current situation. Knowledge about these contexts can then provide predictions useful for further interpreting sensory input and for behaving appropriately in the situation.

In the figure, there are lines leading from the current context representation back to processes that are instrumental in determining what the context is. These lines are meant to indicate that to some extent, the process of context assessment includes circularities. The current context affects what is perceived to be the current situation, which contexts are evoked by the perceived situation, and even the diagnosis of the (next) context.

2 Context-Mediated Behavior

We have developed an approach to context-sensitive reasoning based on the model discussed above. This work was begun over ten years ago in the domain of medical diagnosis [12; 13] and is currently continuing in the domain of controlling autonomous intelligent agents [14; 13; 15; 1], such as autonomous underwater vehicles. We have just begun extending the work to the domain of multi-modal interfaces [11] and to multiagent systems, in particular autonomous oceanographic sampling networks [16].

We call our approach *context-mediated behavior* (CMB) [1] because it is based on the idea that all behavior depends upon, or is mediated by, context. The key ideas in CMB are: contexts should be explicitly represented; contextual knowledge should be associated with context representations; and contextual knowledge should guide all facets of an agent's behavior.

In this section, we describe CMB. We first discuss the knowledge structures that represent contexts, *contextual schemas*, and then how those schemas are used. In the next section, we discuss CMB in relation to several AI systems.

2.1 Contextual Schemas

In our approach, contexts are represented by *contextual schemas* (c-schemas). C-schemas are frame-like knowledge structures that are descendents of Schank's *memory organization packets* (MOPs) [17]. They are organized in a content-addressable conceptual memory similar to the CYRUS [18] program.

C-schemas contain both *descriptive* and *prescriptive* knowledge about the contexts they represent. The descriptive knowledge consists of:

- Features of the situation that must be present (or not present) in order for it to be considered an instance of the context. This allows the agent to "diagnose" the current situation as an instance of a context it knows about.
- Features of the situation, perhaps yet unseen, that are expected in this context. This allows the agent to make predictions about things it is likely to see that may affect problem solving (e.g., to allow it to recognize anticipated events). It also allows the agent to disambiguate sensory input based on contextual, top-down predictions.
- Context-specific ontology/meaning of particular concepts. Concepts often have different meanings in different contexts; contextual knowledge provides this information to the agent. For example, changes in the meaning of fuzzy linguistic values can be handled by storing context-specific membership functions for the values in c-schemas [19]. Similarly, neural networks could be made to recognize different things in different contexts by storing context-specific weights in c-schemas.

Prescriptive knowledge, that is, information about how to behave in the context, is also stored in c-schemas. This includes information about:

- handling unanticipated events: how to detect them, how to diagnose their meaning in the context, their context-specific importance, and how to appropriately handle them;
- focusing attention: which goals should/should not be worked on and how important particular goals are in the context;
- goal-directed behavior: knowledge about how to achieve goals appropriately in the situation;
- non-goal-directed behavior: knowledge governing the expression of behavior that is not directly related to goals, such as turning off obstacle avoidance when an AUV is in the context of docking, etc; and
- new goals that should be pursued because the agent is in the context.

2.2 Context Assessment in CMB

In CMB, a *context manager* module is responsible for context assessment and ensuring that contextual knowledge is distributed to other modules as needed. In the Orca AUV mission controller, the context manager is being implemented as ECHO (Embedded Context-Handling Object) [1].

To do context assessment, the context manager first retrieves the appropriate c-schemas from a schema memory based on features observed in the current situation. This corresponds to the perceived situation evoking candidate contexts in the model above. It then uses information from these c-schemas to do differential diagnosis. The process proposed for differential diagnosis in CMB is that developed by Miller and colleagues and Feltovich for medical differential diagnosis [7; 20]. C-schemas are grouped into logical competitor sets, each of which defines a diagnostic problem to be solved. How strongly the c-schema is evoked by the situation affects its rating. Predictions from a c-schema that are satisfied in the perceived situation increase the confidence that the c-schema fits the situation; violated predictions decrease the confidence. Comparisons between the c-schemas drive the focus of diagnosis.

The product of the diagnostic process is a set of c-schemas, each of which represent a context that describes the situation. These are then merged by the context manager to create a complete picture of the context, called the *context object*. Knowledge from this representation is then used to affect all aspects of the agent's behavior.

The context manager parcels out information from the context object to the rest of the agent. In ECHO, the current plan is to implement this as follows. The agent's other modules register their interests with the context manager, much as agents register with facilitators in multiagent systems based on KQML (knowledge query and manipulation language) [21]. When the a new context has been diagnosed, ECHO will either tell the interested modules that there is a new context or send them the information they requested, depending on how they registered.

The context manager constantly monitors the situation to determine if the context has changed. If so, then a new context object is created. One way that it can notice a changed context is when new c-schemas are evoked based on new features in the situation. It can also explicitly examine the current situation from time to time to decide if it is still adequately described by the context object.

This process is similar to case-based reasoning (CBR) in many respects. Indeed, the process is part of an overall approach called *schema-based reasoning* [13] that grew out of and is a generalization of CBR. Differences from CBR include the use of generalized rather than particular cases and using diagnostic reasoning to select the (general) cases to use.

3 CMB in Intelligent Agents

CMB was partially implemented in a medical diagnostic reasoner, and a full version is currently being implemented in the Orca AUV controller. In addition, we have plans to test the approach in other applications, as discussed here.

MEDIC. MEDIC [12; 13] was a schema-based diagnostic reasoner whose domain was pulmonary infections. It grew out of both work in case-based reasoning and reactive planning. It was an *adaptive, schema-based reasoner*, using generalized cases to control its reasoning and capable of changing the way it behaved based on techniques from reactive planning [22] research.

Contextual information is very important in medical diagnosis. The meaning of a sign (objective finding) or symptom (subjective) depends on context; for example, a persistent cough in a young, generally healthy person should make the diagnostician think of something different than when observed in a chronic smoker or an inner city dweller with HIV (e.g., respiratory infection, cancer, and tuberculosis, respectively).

In MEDIC, the contexts we were interested in had to do with patient presentation. These were early on called *diagnostic MOPs*, then later the name was changed to "contextual schemas" as it became clear that they were an instance of a larger, more generally-useful class of knowledge structures. Each c-schema represented a picture of the current diagnostic session centered around the patient presentation. For example, MEDIC had c-schemas for “consultation”, “cardiopulmonary consultation”, and “cardiopulmonary consultation in which the patient is an alcoholic”. Contextual schemas in MEDIC were monolithic structures representing the entire problem solving context; the best c-schema returned by memory was used as the context object (though not referred to by that name). In a superficial way, MEDIC's c-schemas were similar to earlier work on prototypes in diagnosis by Aikins [23].

Other contexts are important in medicine that MEDIC did not examine. For example, diseases themselves, or rather their presence, define contexts; indeed, the ultimate goal of a purely diagnostic program is to determine the current context to the level of what disease (or set of diseases) is present in the patient. Disease contexts can provide additional information that is very important, such as the prognosis and suggestions for treatment.

MEDIC ignored an important feature of contexts in general, and in medicine in particular: the evolution of contexts over time. For instance, a context defined by patient P having disease D has many “sub-contexts” corresponding to the evolution of the disease, possibly in response to treatment. This fluid nature of contexts is very difficult to capture in AI knowledge structures. MEDIC was concerned with diagnosing “snapshots” of a patient, similar to the clinicopathological conference (CPC) exercises that doctors engage in, or to diagnosis on an outpatient basis. Consequently, tracking the patient through time was not necessary.

There was, however, some evolution of contexts during a session with MEDIC. As the program's understanding of the case grew as findings were presented and questions answered, different c-schemas would match the situation. Usually, the new c-schemas were specializations of the old, allowing fine-tuning of MEDIC's behavior, but sometimes the c-schema would correspond to a different context altogether, which would change the hypotheses MEDIC was considering.

Orca. The CMB process as described above is being implemented in Orca, an intelligent mission controller for oceanographic AUVs [14; 13; 15; 1]. In particular, the ECHO context manager will overcome some of the limitations of MEDIC's approach. It will diagnose c-schemas and merge them into a coherent picture of the overall context. A variety of context types is being considered, for example having to do with the vehicle, the environment, and the mission. Orca will have c-schemas, for example, that represent the contexts “has low power”, “on a search mission”, “in a harbor”, “in Bar Harbor”, and so on. Some work has also begun on handling the changing character of the situation while within a context; for example, the physical properties of the environment change as an AUV transits a harbor, but throughout it makes sense to consider the AUV in the context of “in a harbor”.

We will not further describe ECHO here, as it is adequately described by the preceding section and elsewhere [1]. However, it is instructive to see how a context manager such as ECHO is integrated into an agent. Figure 2 shows the internal structure of the Orca program. As can be seen, ECHO watches the current situation, attempting to detect when the context changes. It detects this in part when new c-schemas are retrievable from the long-term memory based on the changing situation's features. When this happens, it reassesses the context and possibly forms a new context object. Information from the object is then made available to Orca's other modules. Knowledge about event handling is sent to Event Handler. This includes information useful for detecting events as well as diagnosing their cause, assessing their importance, and responding to them. Knowledge about goals is sent to the Agenda Manager. This module uses the knowledge of goal importance and appropriateness to manage Orca's focus of attention. Knowledge about appropriate ways to

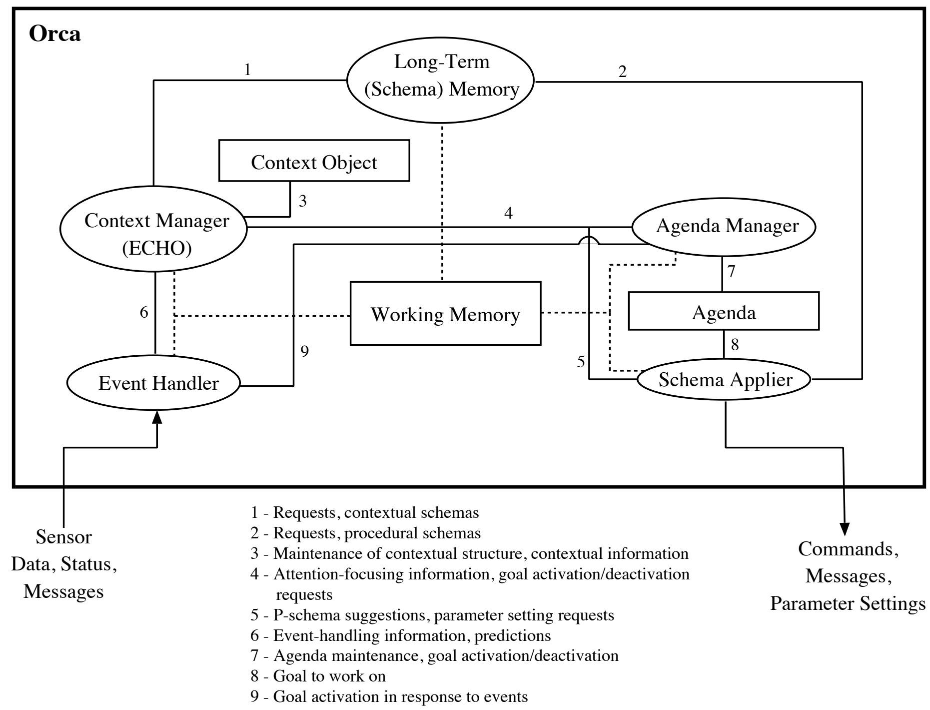

Figure 2: Internal structure of Orca. After Turner [1998].

achieve goals is sent to Schema Applier. This knowledge is in the form of suggestions of procedural schemas (p-schemas) that are likely to be useful for the goals in the current context. (P-schemas are packets of procedural knowledge that can represent plans, scripts, or rules.) Knowledge about the context-dependent meaning of concepts, as well as predictions about unseen features of the world, are available to all modules.

CoDA. We are beginning to look at the role and use of contextual knowledge in multiagent systems. Our initial domain is the control of an autonomous oceanographic sampling network (AOSN) [24], and our project is CoDA (Cooperative Distributed AOSN controller) [16].

We are just beginning to look at context in this setting. Some of the issues that come to mind are how to use context to help select organizational structures, to select communication modes and channels, and to recognize and respond to opportunities to reorganize the system.

Other interesting issues to be explored include the notion of shared context between the agents and how the agents can agree on what the context is. We believe that explicit context representation is critical in this case, since agents must be able to agree on the context they are in. While it will not be necessary (or perhaps even possible in the general case) for the agents to have identical representations of contexts, explicit context representation will allow them to reason about their own contexts as well as others', and to communicate about contexts, in order to come to an agreement about what the current shared context is.

We believe CMB can be extended to the multiagent case, and we plan to explore this in the near future.

Multi-modal interfaces. We have also begun to examine the role of context in multi-modal (natural language and graphics) interfaces to geographical information systems. This work is describe elsewhere [11; 25]. Briefly, we have identified several kinds of contexts (or components of the context) active in this application: the natural language discourse context; the graphics context, including a graphical equivalent to discourse context; the task context; the context defined by the kind of user and the particular user; a context defined by the location(s) being discussed; the temporal context of what is being discussed (e.g., "a building used to be here" versus "there is a building here now"); the context of where the system is being used; and a context defined by explicitly-identified symbol–object mappings. We call each of these kinds of contexts *contextual aspects*, since they describe aspects of the context present in each session. Two of these aspects are similar to four of the "context-space" dimensions recently discussed by Lenat [2].

We are investigating the use of contextual knowledge from representations of these aspects to understand ellipsis and other phenomena important to understanding multi-modal communication. We intend to investigate the applicability of CMB to a multi-modal GIS interface in the near future.

4 Conclusion

Context profoundly affects the appropriateness of an agent's behavior. Consequently, an agent needs to take the context into account when deciding how to behave. In this paper, we have argued that this is best done by the agent having an explicit representation of contexts it may find itself in as well as having contextual knowledge about how to behave in those contexts. We presented a model of how intelligent agents can assess the context they are in and use that to guide their behavior.

We have developed an approach to implementing this model, context-mediated behavior. CMB was partially implemented in a medical diagnostic program and is now being fully implemented and tested in Orca, a schema-based controller for autonomous intelligent agents. In the near future, we will apply our approach to multi-modal interfaces and to multiagent systems.

5 Acknowledgements

The author would like to thank Elise Turner and other members of the University of Maine Cooperative Distributed Problem Solving Research Group as well as Patrick Brézillon for many helpful discussions.

References

1. R. M. Turner. Context-mediated behavior for intelligent agents. *International Journal of Human–Computer Studies*, 48(3):307–330, March 1998.
2. D. Lenat. The dimensions of context-space. Published on-line at URL http://www.cyc.com/context-space.rtf, accessed March 28, 1999., 1998.
3. R. E. Fikes and N. J. Nilsson. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2:189–208, 1971.
4. D. E. Wilkins. Domain-independent planning: Representation and plan generation. *Artificial Intelligence*, 22(3), 1984.
5. R. Davis and B. G. Buchanan. Meta-level knowledge. In B. G. Buchanan and E. H. Shortliffe, editors, *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Addison–Wesley Publishing Company, Reading, Massachusetts, 1984.
6. P. Brézillon, C. Gentile, I. Saker, and M. Secron. SART: A system for supporting operators with contextual knowledge. In *Proceedings of the 1997 International and Interdisciplinary Conference on Modeling and Using Context (CONTEXT–97)*, pages 209–222, 1997.

7. R. A. Miller, H. E. Pople, Jr., and J. D. Myers. INTERNIST–1, an experimental computer-based diagnostic consultant for general internal medicine. *New England Journal of Medicine*, 307:468–476, 1982.

8. D. F. Noble. Schema-based knowledge elicitation for planning and situation assessment aids. *IEEE Transactions on Systems, Man, and Cybernetics*, 19(3):473–482, May/June 1989.

9. E. H. Shortliffe. *Computer-based Medical Consultations: MYCIN.* Elsevier, New York, 1976.

10. P. Öztürk and A. Aamodt. Towards a model of context for case-based diagnostic problem solving. In *Proceedings of the 1997 International and Interdisciplinary Conference on Modeling and Using Context (CONTEXT–97)*, 1997.

11. E. H. Turner, R. M. Turner, J. Phelps, C. Grunden, M. Neale, and J. Mailman. Aspects of context for understanding multi-modal communication. In *Proceedings of the 1999 International and Interdisciplinary Conference on Modeling and Using Context (CONTEXT-99)*, Trento, Italy, 1999.

12. R. M. Turner. When reactive planning is not enough: Using contextual schemas to react appropriately to environmental change. In *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society*, pages 940–947, Detroit, MI, 1989.

13. R. M. Turner. *Adaptive Reasoning for Real-World Problems: A Schema-Based Approach.* Lawrence Erlbaum Associates, Hillsdale, NJ, 1994.

14. R. M. Turner and R. A. G. Stevenson. Orca: An adaptive, context-sensitive reasoner for controlling AUVs. In *Proceedings of the 7th International Symposium on Unmanned Untethered Submersible Technology (UUST '91)*, pages 423–432, 1991.

15. R. M. Turner. Intelligent control of autonomous underwater vehicles: The Orca project. In *Proceedings of the 1995 IEEE International Conference on Systems, Man, and Cybernetics.* Vancouver, Canada, 1995.

16. R. M. Turner and E. H. Turner. Organization and reorganization of autonomous oceanographic sampling networks. In *Proceedings of the 1998 IEEE International Conference on Robotics and Automation (ICRA'98)*, pages 2060–2067, Leuven, Belgium, May 1998.

17. R. C. Schank. *Dynamic Memory.* Cambridge University Press, New York, 1982.

18. J. L. Kolodner. *Retrieval and Organizational Strategies in Conceptual Memory.* Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1984.

19. R. M. Turner. Determining the context-dependent meaning of fuzzy subsets. In *Proceedings of the 1997 International and Interdisciplinary Conference on Modeling and Using Context (CONTEXT-97)*, Rio de Janeiro, 1997.

20. P. J. Feltovich, P. E. Johnson, J. A. Moller, and D. B. Swanson. LCS: The role and development of medical knowledge and diagnostic expertise. In W. J. Clancey and E. H. Shortliffe, editors, *Readings in Medical Artificial Intelligence*, pages 275–319. Addison–Wesley Publishing Company, Reading, Massachusetts, 1984.

21. R. Patil, R. Fikes, P. Patel-Schneider, D. McKay, T. Finin, T. Gruber, and R. Neches. The DARPA Knowledge Sharing Effort: Progress report. In *Proceedings of the Third International Conference on Principles of Knowledge Representation and Reasoning*, 1992.

22. M. P. Georgeff and A. L. Lansky. Reactive reasoning and planning: An experiment with a mobile robot. In *Proceedings of the Sixth National Conference on Artificial Intelligence*, pages 677–682, Seattle, Washington, 1987.

23. J. S. Aikins. *Prototypes and Production Rules: A Knowledge Representation for Computer Consultations.* PhD thesis, Stanford University, 1980.

24. T. Curtin, J. Bellingham, J. Catipovic, and D. Webb. Autonomous oceanographic sampling networks. *Oceanography*, 6(3), 1993.

25. E. H. Turner, R. M. Turner, C. Grunden, J. Mailman, M. Neale, and J. Phelps. The need for context in multi-modal interfaces. In *Workshop Notes for the 1999 AAAI Workshop on Reasoning in Context for AI Applications*, AAAI Technical Report (ISBN 1-57735-098-7), pages 91–95, Orlando, FL, July 1999. AAAI Press.