

# Automatic Development of an Abstract Context Model for an Intelligent Environment

Oliver Brdiczka, Patrick Reignier, and James L. Crowley

*Laboratoire GRAVIR, INRIA Rhône-Alpes, France  
{brdiczka, reignier, crowley}@inrialpes.fr*

*Abstract*—This paper addresses the problem of learning in intelligent environments. An intelligent environment perceives user activity and offers a number of services according to the perceived information about the user. An abstract context model in the form of a situation network is used to represent the intelligent environment, its occupants and their activities. The objective is to adapt the system services, which are associated to the situations of the model, to the changing needs of the user. For this, a supervisor gives feedback by correcting system services that are found to be inappropriate to user needs. The situation network can be developed by exchanging the system service-situation association or by *splitting the situation*. The situation split is interpreted as a replacement of the former situation by sub-situations whose number and characteristics are determined using conceptual or decision tree algorithms. Different algorithms have been tested on a context model within the SmartOffice environment of the PRIMA research group. The decision tree algorithm (ID3) has been found to give the best results.

*Index Terms*—situation split, context model learning

## I. INTRODUCTION

An environment is called “perceptive” if it is capable of maintaining a model of its occupants and their activities. An environment becomes “active” when it is capable of (re)actions (system services). An “interactive” environment is based on the capacity of perception, (re)action and communication with the users. “Intelligent” environments should provide services while minimizing disruptions, such as explicit man-machine communication. This requires that the intelligent environment perceives user activities and identifies user needs correctly in order to react in an appropriate way. However, user needs and system services evolve in the course of time. Further, different users may have different needs for the same activities. Thus an intelligent environment must be capable of adapting and developing its services automatically to meet the specific needs of a user.

In this paper, we describe a method for evolving and developing the abstract context model for an intelligent environment. Such evolutions are intended to maximize the correctness of the executed system services concerning the user needs perceived by a supervisor.

## II. PROBLEM STATEMENT

The problem addressed in this paper is machine learning in an intelligent environment. The intelligent environment is a computer system that executes a number of services according to perceptual information on user actions or activity. As we know, user behavior changes in the course of time. The automatic adaptation of system (re)actions according to changing user needs is seen as machine learning process. We need information in the form of feedback on executed system services in order to guide the learning process. Further, we can define several qualities that this machine learning process should have:

### A. Understandable Representation and Reasoning

We consider that a user is only willing to accept an intelligent environment offering services implicitly if he understands and foresees its decisions. Thus we want the user to be able to understand the system perceptions and their representation in the model of the interactive environment. Further, the learning process, i.e. the reasoning and the development of the model necessary to cover changing user needs, should also be understandable for the user.

### B. Supervisor corrections (feedback)

We want to minimize the frequency with which the system offers inappropriate services, while minimizing disruption. This means that the feedback given to the system is to be minimal to achieve the wanted changes of system services. We assume that a person, denoted *supervisor* in the following, is capable of specifying system services to be executed by the system and that his feedback is always consistent. The user himself or another person can act as this supervisor.

In this paper, we distinguish three forms of supervisor feedback:

- (Re)action correction: the service or (re)action executed by the system is wrong and a different service must be executed instead. The supervisor gives the different system (re)action as feedback to the system. This includes the case where the supervisor wants the system to execute a (re)action while the system does not execute anything.
- (Re)action deletion: the service executed by the system is wrong and no system service must be executed

instead. The supervisor gives a particular (re)action, the “erase” (re)action, as feedback to the system.

- (Re)action preservation: the service executed by the system is correct. The supervisor does not give any information to the system. As we assume that the supervisor is always consistent, we can interpret the absence of his corrections or deletions as positive feedback for the currently executed system (re)actions.

The supervisor will not give any further feedback on the model representing the intelligent environment. The learning process must integrate the information given on the system (re)actions by evolving and developing the model.

### III. PREVIOUS APPROACHES AND RELATED WORK

An early example of an interactive environment is the KidsRoom [1], a perceptually-based, interactive, narrative playspace for children. The KidsRoom environment interacts with the children in order to narrate a story. It does not have, however, any learning capacity concerning the automatic development of its (re)actions. The perception modules and the (re)actions in the environment are integrated in a preprogrammed story context and need to be adapted by hand (in general by a programmer).

The ContAct project realized a system [2] deriving user activity and availability from sensor data in an office environment. The system uses a naive Bayesian Classifier to learn user activity and availability directly from sensor data according to given user feedback. This system is, however, only perceptive and not (inter)active as there are no (re)actions defined or planned by the system. Further, no understandable representation of the environment is provided.

The Lumiere project [3] at Microsoft Research realized methods and an architecture for reasoning about the goals and needs of software users as they work with software. The objective is to learn appropriate interactions with the user according to perceived user activities in order to offer assistance within a software environment. At the heart of Lumiere are Bayesian models, which are based on a large amount of example data from users and experts. Although these models could identify a number of (understandable) pertinent variables of human behavior (like task history or assistance history), their reasoning is not obvious to the user.

In this paper, we want to focus on an understandable model of the environment whose initial construction can be done by hand (without large amounts of example data). The context model [4] is a non-Bayesian model inspired by concepts of planning and knowledge representation used in robotics and artificial intelligence. The central concept of this model is the notion of situation. A situation refers to a particular state of the environment. It is composed of a particular configuration of entities, roles and relations (see Fig. 1 for an example).

An entity may generally be understood as corresponding to a physical object or person. It is created accompanied by a set

of numerical or symbolic properties. We say that an entity is observed to play a role if it passes a role acceptance test on its properties. This role acceptance test may be seen as a predicate function defined over entities and their properties. A person may, for example, play the role *lecturer* if he or she stands next to the presentation screen, which is tested by comparing the person's position to predefined values (acceptance test).

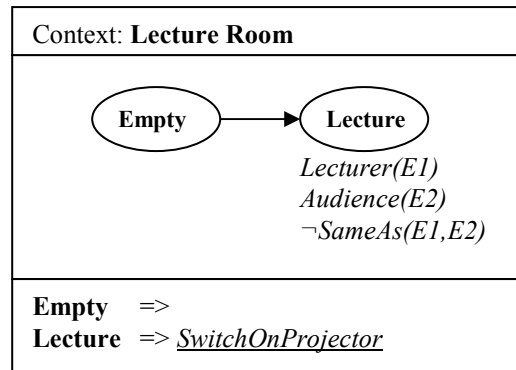


Fig. 1. Example of a context model for a lecture room. **Empty** and **Lecture** are the available situations. *Lecturer*, *Audience* are the available roles and *SameAs* the available relation. *SwitchOnProjector* is a system (re)action. *E1* and *E2* are entities.

A relation is defined as a predicate function on several entities and their properties. The identity relation may, for example, be created by comparing the names of two entities. A context is determined by the available roles and relations. A situation stands for a particular assignment of entities to roles completed by a set of relations between entities. Thus a change in the relations between entities or in the bindings of entities to roles results in a change of situation. This (possible) change of situation is represented by an arc connecting these situations. The context can then be represented by a network of situations. System (re)actions (system services) are directly associated to the situations in the network.

### IV. METHOD

The understandable form of the context model allows an easy predefinition of situations and associated system (re)actions by a supervisor. The learning process must adapt the predefined context model according to the given supervisor feedback on the system (re)actions. As the different layers of a context model (entities - roles, relations - situations) influence each other, the learning process cannot adapt them simultaneously. Thus we will focus on the development of the situation network and the associated system (re)actions.

Bayesian models (in particular Hidden Markov Models [5]) as well as algorithms based on first-order logic [6] have been considered to represent and adapt the situation network. However, these approaches do not have good properties concerning the extension of the number of situations, which is essential for developing a situation network. Bayesian models require a large amount of example data to extend the number

of states. First-order logic algorithms cannot create new predicates (problem of higher order logic), which corresponds to the extension of the number of situations. Thus the method proposed in this paper is based on algorithmic changes of the structure of the situation network.

### A. Algorithm

Fig. 2 shows an overview of the proposed algorithm. The input of the algorithm is a predefined situation network (context model) and feedback given by a supervisor. The supervisor corrects, deletes or preserves the (re)actions executed by the system while observing a user in the environment. Each correction, deletion, or preservation generates a training example for the learning algorithm containing current situation, roles and relations configuration, and the (correct) (re)action. The differences between the (re)actions given in the training examples and the (re)actions provided in the predefined situation network will drive the different steps of the algorithm.

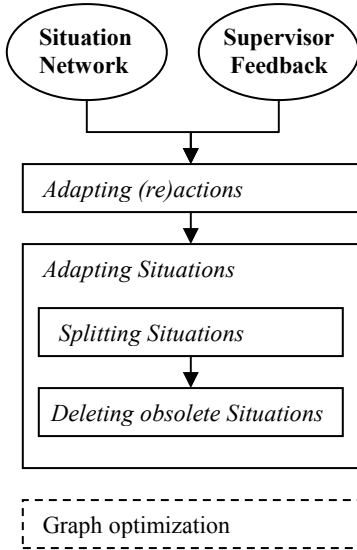


Fig. 2. Overview of the different steps of the algorithm

In the first step, the algorithm tries to directly modify its (re)actions using the existing situation network. If (re)action A is associated to situation S, and all training examples indicate that (re)action B must be executed instead of A, then B is associated to S and the association between A and S is deleted.

In the second step, the algorithm tries to modify the situation network. The situation split is executed when the supervisor perceives several situations (expressed by different (re)actions in the training examples) while the predefined situation network only perceives one situation (expressed by one (re)action). Thus the situation perceived by the predefined situation network may be too general and the algorithm tries to split it in sub-situations. The role, relations configuration of these sub-situations needs to be determined according to the given training examples (Section B). After the situation split, (sub)situations whose associated (re)action is the “erase” (re)action are deleted.

The last step is the optimization of the resulting situation network. This is seen as optional step (dotted line) and will not be prioritized in this paper. The elimination of obsolete arcs or the detection of obsolete structures in the network is conceptually possible, but not demonstrated by the implementation described here.

### B. Splitting Situations

When splitting situations, a number of training examples indicate different (re)actions for one situation of the predefined situation network. Several sub-situations need to be created for these (re)actions. We must determine the characteristic role, relation configurations of these sub-situations (Fig 3).

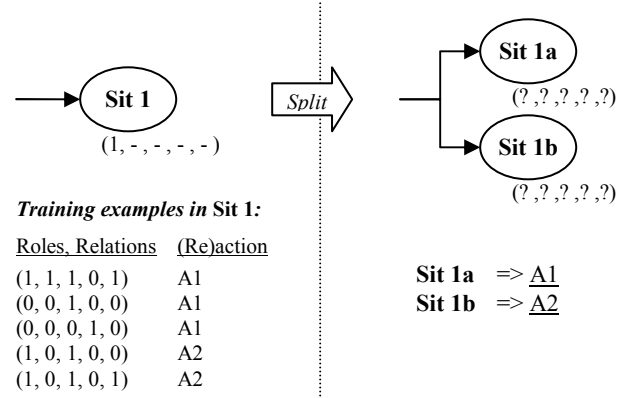


Fig. 3. Splitting Situations. The characteristic roles, relations configuration of **Sit 1a** and **Sit 1b** need to be determined (“?”).

To facilitate the processing of roles and relations, they are represented as a fixed-sized vector containing one 0/1 value for each available role and several 0/1 values for each available relation. The value 1 means that the corresponding role or relation is valid; the value 0 means that the role or relation is not valid. A characteristic role, relation configuration for one situation may contains blanks (“-”) for those roles or relations that are not characteristic for this situation. A training example contains a vector with specific values reflecting the current roles, relations configuration when recording the training example and the corresponding (re)action (given by the supervisor).

The determination of the characteristic roles, relations configurations of the sub-situations can be seen as classification problem. The (re)action labels of the training examples can be interpreted as class labels. For each class, we need then to determine the concepts or hypotheses based on the given role, relation vectors of the class. These concepts or hypotheses are then used to construct the characteristic role, relation configurations of the corresponding sub-situation.

The first considered learning method is the conceptual learning algorithm Find-S ([7] chapter 2.4) which constructs the most specific hypothesis for each (re)action based on the role, relation configurations in the given training examples (Fig. 4). The resulting hypotheses for the created sub-situations often contain, however, specific values for the existence or non-existence of roles or relations that are not necessary or characteristic. As a consequence, small variations

in the role, relation configuration may not be covered by the created sub-situations because their hypotheses are too specific.

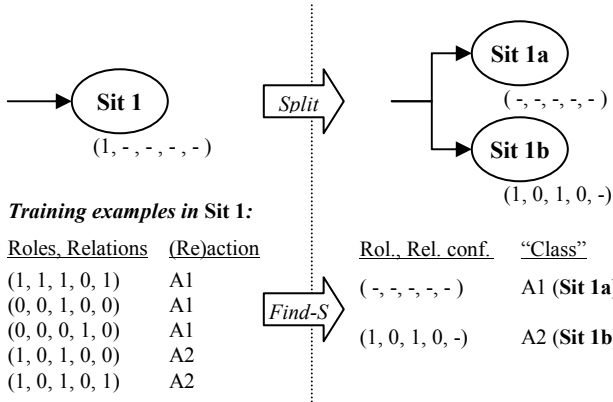


Fig. 4. Splitting Situations with Find-S algorithm.

To produce more general hypotheses for the sub-situations, we consider the conceptual learning algorithm Candidate Elimination ([7] chapter 2.5). This algorithm constructs the most specific and the most general hypotheses for each (re)action based on the role, relation configurations in the given training examples. By combining the most general hypotheses for each (re)action, we construct the role, relation configuration for the corresponding sub-situations (Fig. 5).

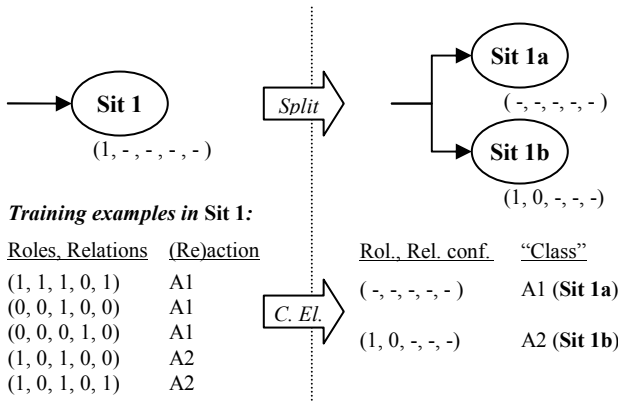


Fig. 5. Splitting Situations with Candidate Elimination algorithm.

Both algorithms Find-S and Candidate Elimination have, however, the restriction that they can only find one conjunctive concept for each (re)action, i.e. if the training examples indicate that a (re)action is to be executed in two different complementary role, relation configurations, Find-S and Candidate Elimination will fail to construct several hypotheses (and thus sub-situations) for this one (re)action. This is due to the fact that neither algorithm can construct disjunctive hypotheses.

We consider Decision Tree learning methods, in particular the algorithm ID3 [8], in order to address the limitation of conceptual learning methods. The idea is to construct a decision tree that classifies the different (re)actions found in the training examples of one situation (Fig. 6).

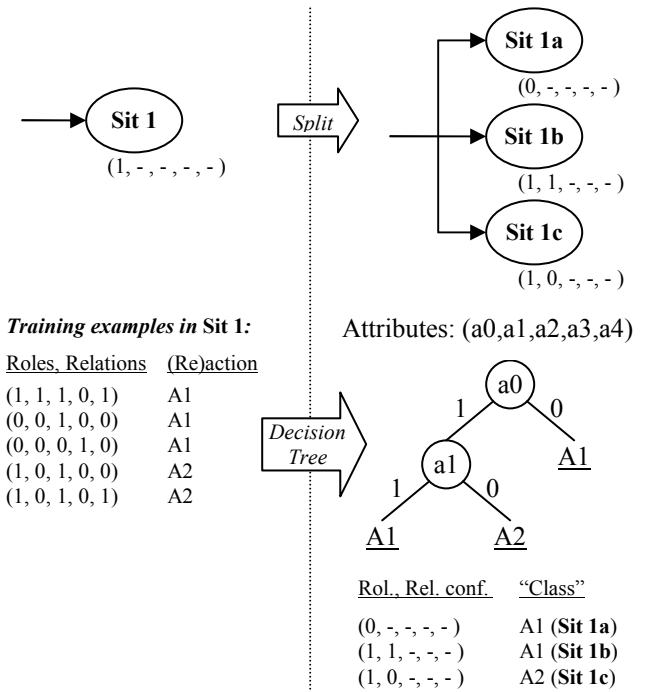


Fig. 6. Splitting Situations with Decision Tree algorithm.

The attributes of this decision tree are the roles and relation values (0/1 values of the vector). Each leaf of the tree is labeled with a (re)action (class). The path from the root of the tree to the leaf gives the characteristic role, relation configuration for the sub-situation to be created for this (re)action. We can have several leaves with the same (re)action, which corresponds to the creation of several sub-situations for this (re)action (disjunctive hypotheses).

## V. IMPLEMENTATION

A context model for office activity within the SmartOffice environment [9] of the PRIMA group has been designed and implemented. In this environment, entities are created by a robust tracking system [10]. The position of the created entities determines several roles like *comes\_in* or *works\_on\_PC*. Additional roles are determined by the login of an entity (person) to a computer in the environment or specific appointments marked in the agenda of the logged entity (person). The *not\_same\_entity\_as* relation is used to distinguish entities in the environment. The (re)actions of the system are based on the control of the Linux music player and the projection of different messages or presentations on different surfaces in the environment. The learning algorithms run on data base tables containing a representation of the current situation network and the training examples. A control process programmed in the forward chaining rule programming environment Jess [11] is used to execute the situation network. This situation network represented by rules is automatically generated from the data base tables of the learning algorithms. The supervisor feedback can't be given while the user is acting in the environment (i.e. while the

control process is running). Thus the control process and the learning algorithms need, at present, to run sequentially and not in parallel.

## VI. EVALUATION AND RESULTS

To evaluate our method, two experiments have been executed on the predefined context model of the SmartOffice environment. Both experiments have the same goal concerning the evolution of the system services. The supervisor gives feedback based on this goal during both experiments. We do a cross-validation by adapting the predefined situation network using the supervisor feedback of the first experiment and by evaluating the second experiment on the adapted situation network (and inverse). The evaluation is done on the number of correctly classified training examples, i.e. correctly executed (re)actions, as well as on the review of the adaptations of the predefined situation network.

<i>Find-S</i>	<b>A0</b>	<b>A8</b>	<b>A9</b>
<b>A0</b>	0.87	0.04	0.09
<b>A8</b>	0.50	0.50	0.00
<b>A9</b>	0.50	0.00	0.50

Table 1. Confusion matrix for (re)action execution (Find-S).

The goal of both experiments was to adapt the situation network in order to cover the correct turn-on and turn-off of the Linux music player depending on the activities (=roles, relations) of the user. Table 1, 2 and 3 show the results of the (re)action execution in the form of confusion matrices. (**A8** switches on the music player, **A9** switches off the music player, and **A0** is the “do nothing” (re)action).

<i>C. El.</i>	<b>A0</b>	<b>A8</b>	<b>A9</b>
<b>A0</b>	0.91	0.04	0.04
<b>A8</b>	0.66	0.33	0.00
<b>A9</b>	0.75	0.00	0.25

Table 2. Confusion matrix for (re)action execution (Candidate Elimination).

In all experiments, the structural development of the situation network corresponds to the expected changes. Concerning the correct classification of the training examples, i.e. the correct execution of the (re)actions, the Decision Tree algorithm (ID3) gives the best results.

<i>D. Tr.</i>	<b>A0</b>	<b>A8</b>	<b>A9</b>
<b>A0</b>	0.83	0.09	0.09
<b>A8</b>	0.00	1.00	0.00
<b>A9</b>	0.00	0.00	1.00

Table 3. Confusion matrix for (re)action execution (Decision Tree).

The improved results of Decision Tree approach are due to the fact that this algorithm supports disjunctive hypotheses. However, the Decision Tree algorithm tends to construct “too general” hypotheses for the sub-situations, which can lead to several inappropriate classifications.

## VII. CONCLUSION

We have presented a learning method for evolving system services to changing user needs in an intelligent environment. The intelligent environment has been modeled as a situation network. This network is adapted according to feedback given by a supervisor using an algorithmic learning method.

The results of the method are encouraging. The system services desired by the human supervisor are correctly integrated into the situation network structure. However, the method relies on the correct detection of roles and relations, which are seen as perception modules encapsulating perception error handling. Further, the supervisor feedback needs to be consistent, which is not always the case in reality. Thus the focus of our future research will concern the extension of the context model and of the learning algorithms to fuzzy or probabilistic values by integrating the confidence values of the perceptual processes.

## REFERENCES

- [1] A.F. Bobick, S.S. Intille, J.W. Davis, F.Baird, C.S. Pinhanez, L.W. Campbell, Y.A. Ivanov, A. Schutte, and A. Wilson. The KidsRoom: A Perceptually-Based Interactive and Immersive Story Environment. In *Presence (USA)*, 8(4):369-393, August 1999
- [2] M. Muehlenbrock, O. Brdiczka, D. Snowdon, and J.-L. Meunier. Learning to Detect User Activity and Availability from a Variety of Sensor Data. In *IEEE International Conference on Pervasive Computing and Communications (PerCom '04)*, pages 13-23, March 2004.
- [3] E. Horvitz, J. Breese, D. Heckerman, D. Hovel, and K. Rommelse. The Lumiere Project: Bayesian User Modeling for Inferring Goals and Needs of Software Users. In G.F. Cooper and S. Moral, editors, *Uncertainty in Artificial Intelligence. Proceedings of the Fourteenth Conference (1998)*, pages 256-265, San Francisco, CA, USA, 1998. Morgan Kaufmann Publishers Inc.
- [4] J.L. Crowley, J. Coutaz, G. Rey, and P. Reignier. Perceptual Components for Context Aware Computing. In G. Boriello and L.E. Holmquist, editors, *UbiComp 2002: Ubiquitous Computing. 4<sup>th</sup> International Conference. Proceedings (Lecture Notes in Computer Science)*, volume 2498, pages 117-134, Berlin, Germany, September/October 2002. Springer-Verlag.
- [5] L. R. Rabiner. A Tutorial on Hidden Markov Models and selected Applications in Speech Recognition. *Readings in speech recognition*, pages 267-296, 1990.
- [6] J. R. Quinlan. Learning Logical Definitions from Relations. *Machine Learning*, 5(3):239-266, August 1990.
- [7] T.M. Mitchell. *Machine Learning*. McGraw Hill, New York, USA, international edition, 1997.
- [8] J.R. Quinlan. Induction of Decision Trees. *Machine Learning*, 1(1), 81-106.
- [9] Ch. Le Gal, J. Martin, A. Lux, and J.L. Crowley. SmartOffice: Design of an Intelligent Environment. *IEEE Intelligent Systems*, 16(4):60-66, July/August 2001
- [10] A. Caporossi, D. Hall, P. Reignier, and J.L. Crowley. Robust Visual Tracking from Dynamic Control of Processing. In *Sixth IEEE International Workshop on Performance Evaluation of Tracking and Surveillance*, Prague, Czech Republic, May 2004.
- [11] Jess: the rule engine for Java. <http://herzberg.ca.sandia.gov/jess/>.