From the real-world to its qualitative representation – Practical lessons learned

Gerald Steinbauer and **Jörg Weber** and **Franz Wotawa***
Technische Universität Graz
Institute for Software Technology
8010 Graz, Inffeldgasse 16b/2, Austria
{steinbauer,jweber,wotawa}@ist.tugraz.at

Abstract

In this paper we discuss problems related to extracting qualitative knowledge from sensor inputs of the real physical world. We assume that the required knowledge is represented as facts that can be either true or false in a certain situation. Assigning truth values based on perceptions causes problems because of unreliable sensory inputs or interactions of objects in the real world. We address these problems by using the concept of predicate hysteresis and present experimental results when using hysteresis in robotics. The results indicate an improvement of the overall assignment of truth values.

1 Introduction

Consider a robot that has to provide a certain task at a certain point in time. This robot has to have a knowledge about the physical world not only in terms of quantitative measurements like probability distributions of its location but also in terms of qualitative facts like predicates stating that a ball is in reach. This qualitative representation is necessary for computing actions in order to fulfill the task. Of course this picture of an autonomous agent assumes a symbolic reasoning engine on top which is used to handle high-level control in contrast to low-level control structures which can be implemented as reactive systems.

At the first sight the mapping of quantitative information to its qualitative representation seems not to be big deal and in some cases this is true. For example, when dealing with control systems for a plant with a limited (and known) number of possible interactions with the environment, the mapping problem can be solved by applying the right thresholds and filters. However, in applications like the robotic domain with unpredictable interactions between the robot and its environment the situation changes. Consider for example changes in the light condition. These changes have a substantially impact on the visual perception of the robot. Hence, an object which was within a 1 meter distance before changing the density of light maybe perceived at a higher distance afterwards although the real situation of the relationship between the object and the robot has not changed. Hence, the robot changes its internal state and may choose different actions. A more severe situation can happen when environmental changes cause the robot to switch between two contradicting states, e.g., object in reach and out of reach, which prevents the robot from taking meaningful actions.

The qualitative mapping problem as described before is mainly caused by unreliable perception. Hence, one solution would be to improve the perception algorithms, e.g., the computer vision system, to make it less sensitive to changes of environmental parameters like light conditions. However, the mapping problem itself will not be solved. For example assume a perfect perception system and two robots playing soccer. The situation starts and our robot assumes the ball in reach but the opponent robot kicks the ball slightly. Hence, the situation changes and the ball is no longer in reach. Because of the underlying definition of in reach this might mean that the ball is now more than 1 meter apart from our robot which is also the case for a distance of let us say 1 meter and 1 centimeter. In both situation we would expect our robot to take the same actions but because of a difference of 1 centimeter and the use of a sharp boundary the perceived world is different and thus the actions as well. A solution for this problem could be the introduction of new landmarks. In our example, a new predicate for almost reachable can be introduced. This kind of solution can be compared with solutions for the problem of finding the right qualitative reasoning model for a certain task. [Sachenbacher and Struss, 2001] propose such a solution.

Although, the qualitative mapping problem can be theoretically solved by using perfect sensors and qualitative modeling techniques there is still a need for a practical solution. Perfect sensor input is not available and there is no indication that this problem will be solved soon. This holds especially for visioning systems. Hence, there is a requirement to overcome the problem. In this paper we follow the hysteresis approach from [Fraser *et al.*, 2004]. We introduce the problem again, present a

*Authors are listed in alphabetical order.

practical solution in terms of predicate hysteresis, Discuss experimental results and open issues. The experiments indicate that the use of hysteresis really improves the overall behavior.

2 Symbol Grounding and Action Selection

In order to create a continuous model of the real environment, the perceptions from different sensors (e.g. camera, odometry) are fused. The resulting model contains the positions of objects on the field: the ball, the two goals, and the players. There are different methods for creating world models in dynamic and nondeterministic environments; we use a Kalman filter [Maybeck, 1990] for predictions of object positions and sensor fusion.

This purely quantitative model is transformed to an abstract world model, the *knowledge base*, which is expressed by means of a set K of ground predicates. The knowledge base, which is basically a conjunction of ground atoms, is the source for the qualitative reasoning which is performed by the *Planner* [Fraser, 2003]. The Planner is the strategy layer of the control software for our soccer robots, and its main responsibility is the selection of actions which shall be executed next. The Planner makes use of classic AI planning for creating plans at runtime. It is based on the STRIPS representation language [Fikes and Nilsson, 1972].

This approach has, compared to reasoning based on continuous data, many advantages. Among others, a qualitative model has only a finite number of possible states, and qualitative models are able to cope with uncertain and incomplete knowledge. Another reason is the fact that the programming of the robot is simplified and can also be done by human operators who have no programming skills. The knowledge and the strategy can be neatly expressed in logical formulas.

As already explained, the knowledge of the robot is expressed using ground predicates. The interpretation of a n-ary predicate $p \in P$ bases on the continuous world model M. It can be formalized as follows:

$$I(p(O^n), M) = \begin{cases} true & \text{if } COND_p(O^n, M) = true \\ false & \text{otherwise} \end{cases}$$

A constant O denotes an object of the environment, e.g. *Ball*, *OwnGoal*, or players. The function $COND_p$ is specific for each predicate p. For example, the predicate $inReach(O, M)$ is defined as follows (R is the robot itself):

```
COND_inReach(O, M): boolean
    return (dist(R, O) < 1200)
```

inReach is an example for a predicate whose truth value is grounded on the distance between the robot and another object. For convenience, this kind of predicate is called *distance predicate* from now on. Of course, predicates can state various kinds of knowledge about the environment, for example the visibility of objects (e.g. $unknown(Ball)$ is *true* iff the position of the ball is unknown) or angles between objects.

Based on the current state of the knowledge base, the Planner selects a plan which shall be executed next. A plan P comprises:

1. A precondition PRE_P which is a conjunction of ground literals. P can be executed only if PRE_P is fulfilled.
2. A sequence of actions $< a_1, .., a_n >$. $a_i \in A$ where A is the set of actions the robot is able to execute. A plan is successfully finished iff all actions are finished. The sequence of actions is either dynamically computed using classic AI planning or it is statically defined by a human operator.
3. An invariant INV_P which is a conjunction of ground literals. If an invariant of a currently executed plan is violated, then the plan is aborted.

A more detailed discussion of the plan execution is found in [Fraser *et al.*, 2005].

3 A predicate hysteresis

The mapping from a quantitative model to symbolic predicates in a dynamic and uncertain environment leads to two major problems: First, the truth value of predicates is calculated using thresholds, i.e. there are sharp boundaries. Thus slight changes of the environment can cause truth value changes and result in abortion of plans due to a violation of the invariant, even if the plan still could be finished successfully. The consequence is instability in the high-level decision making process. A *commitment* to a plan, once it is chosen, is desired. Second, sensor data is inherently noisy. Hence, due to the sharp boundaries, sensor noise leads to *unstable knowledge*, i.e. to undesired oscillation of truth values, even if the environment does not change.

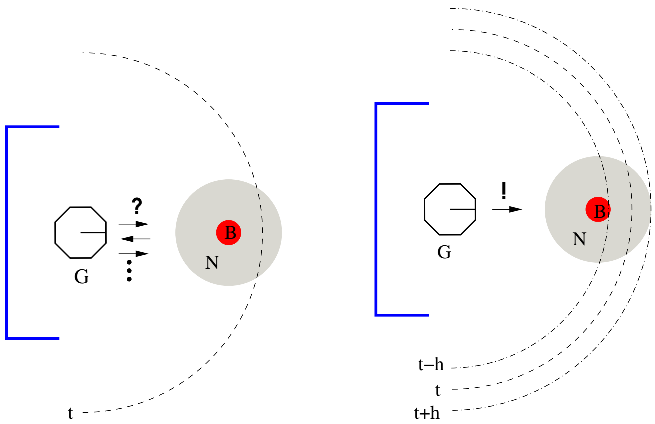

Figure 1: Example: (a) no hysteresis, (b) with hysteresis of size h. G is the goalkeeper, B the ball, the area N depicts the uncertainty of the ball position measurements.

We propose a *predicate hysteresis* as an attempt to mitigate the problems described above. The term *hysteresis* is well known from electrical engineering. It means that the current state is influenced by a decision which has been made previously. We adapt this concept in order to improve the robustness of the decision making process. The basic idea is that, once a predicate evaluates to a certain truth value, only significant changes of the environment can cause a change of this truth value.

Thus an extended interpretation function I_H is introduced:

$$I(p(O^n), M, l) = \begin{cases} true & \text{if } COND_p(O^n, M, l) = true \\ false & \text{otherwise} \end{cases}$$

The variable l represents the current truth value of p. The functions $COND_p$ are also redefined. The general definition for distance predicates is:

```
COND_distance_pred(O, M, l): boolean
 if l then
     return (dist(R, O) < threshold + h
 else
     return (dist(R, O) < threshold - h
```

threshold is specific for each predicate. In the example given above, the predicate *inReach* has a threshold of 1.20m. h is the *hysteresis size*. In this definition, the hysteresis size is defined as an absolute number. In practice, predicates with a larger threshold may also demand a larger hysteresis because in general the sensor noise is higher for distant objects. Thus it is often more convenient to define a hysteresis size h_{rel} as a percentage of the threshold:

```
COND_distance_pred(O, M, l): boolean
 if l then
     return (dist(R, O) < threshold * (1+h_rel)
 else
     return (dist(R, O) < threshold * (1-h_rel)
```

However, in this paper the term *hysteresis size* always denotes an absolute number in mm.

Figure 1 gives an example for the effect of using a predicate hysteresis. It shows a goalkeeper in his goal. The ball has approached the goalkeeper and stops at the position which is shown in the figure. $dist(R, Ball)$ is slightly less than $t - h$, whereas t is the threshold of the predicate *inReach* and h the hysteresis size.

Suppose the goal-keeper's strategy includes the following plans:

	precondition $\equiv$ invariant:	action:
P_1:	$\neg hasBall() \wedge \neg inReach(Ball)$	stay in goal
P_2:	$\neg hasBall() \wedge inReach(Ball)$	grab ball

In (a) as well as in (b), $inReach(Ball)$ becomes *true* and thus P_2 is activated. The goalkeeper starts moving towards the ball.

But in (a), where no hysteresis is used, it may happen that $inReach(Ball)$ becomes $false$ again due to sensor noise. In this case, the current plan is aborted and P_1 is reactivated. This scenario can happen several times in quick succession, the goalkeeper activates plans and aborts them before they can succeed.

In (b) a hysteresis of size h is used. As soon as $inReach(Ball)$ becomes *true* (i.e. the distance between the robot and the ball is less than $t - h$), it keeps this truth value as long as the $dist(R, Ball)$ is less than $t + h$. Thus the truth value of $inReach(Ball)$ is, to a certain extent, robust against the noise of the ball position measurement. In this example, the hysteresis size is sufficiently large to compensate the noise, and the goalkeeper does not abort P_2 after he has made this decision. The goalkeeper *commits* himself to this decision - as it would happen in a real soccer match: if a real goalkeeper leaves his goal in order to grab the ball, he does not change his mind only because of a slight change of the ball position.

4 Experimental Results

The proposed symbol grounding with hysteresis was evaluated on our real robots within the robotic soccer domain. We investigated how the use of a hysteresis in symbol grounding stabilizes the evaluation of the truth value of predicates and reduces the number of undesired changes of the truth value caused by noise and changes in the environment.

We conducted several static and dynamic experiments in which the robot measured the distance to objects on the field using its vision system. Based on these measurements the symbol grounding evaluates the truth value of the distance predicate *inReach*. The distance measurements are not reliable and vary within certain boundaries because of noise and changes in the environmental conditions. Therefore, there are undesired changes in the truth value of predicates even if the distances do not change in the real world.

Figure 2 shows series of distance measurements during a static experiment. The robot was placed 4800 mm away from the yellow goal. We recorded series of distance measurements over periods of 30 seconds. These series were recorded at different times during the day to investigate the influence of changing lighting conditions. Please note that the vision system was calibrated the day before and no adaptation of the vision and camera took place between the different series.

As the experiment setup was totally static a perfect vision system would always report the same distance and there would no change in the truth value of a distance predicate. But the Figure shows that in practice the measurements are affected by noise. Furthermore, it shows the clear dependency of the amount of noise in the data on changing lighting conditions. The extent of noise differs within the different series recorded under different lighting conditions. This change is caused by the fact that the color of objects is differently perceived under changing light and the robot vision relies on the colors of objects. The worst conditions were at 17:00 where it became dark.

Because of the quality of the distance measurements, the symbol grounding with a fixed threshold reports a number of truth value changes of the distance predicate. These changes are undesired because the object positions

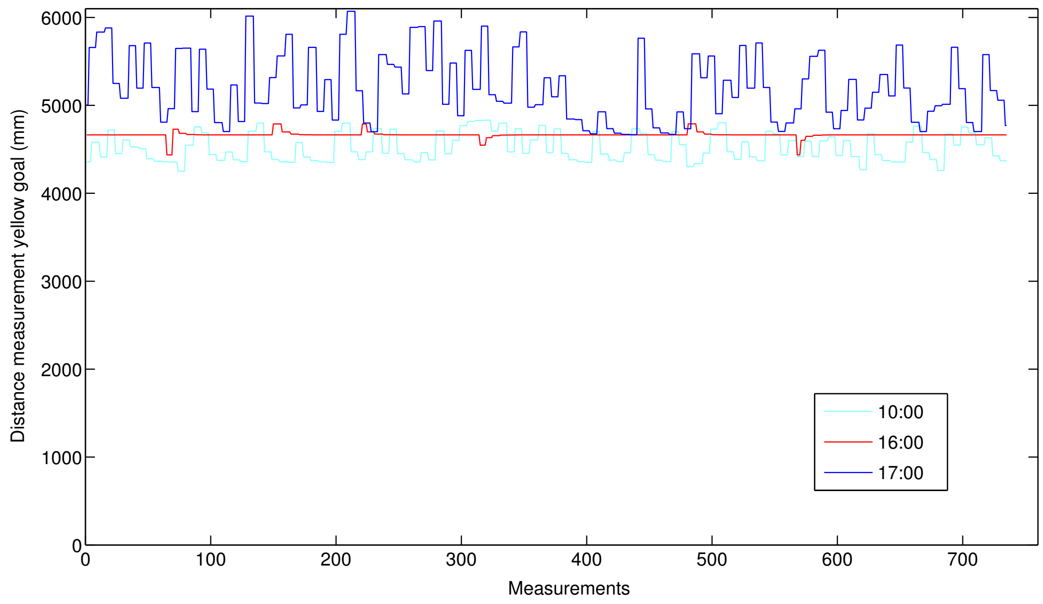

Figure 2: Distance measurements for a static object 4800 mm away from the robot at different times during a day.

# Meas.	μ	σ	Δ	n			
				h=0	h=σ	h=$\frac{\Delta}{3}$	h=$\frac{\Delta}{2}$
	mm	mm	mm				
735	5202	385	1403	61	17	11	0

Table 1: Number of undesired truth value changes n of predicate *inReach* for the yellow goal for static distance measurements at 17:00 with different sizes h for the hysteresis.

did not change in the real world. Table 1 shows the results of the symbol grounding of the series at 17:00 and reports how different sizes of a hysteresis stabilize the symbol grounding. The series contained 735 measurements with a distance mean of 5202 mm and a standard deviation σ of 385 mm. The value Δ is the difference between the maximum and the minimum of the measured distance within this series. If we do not use a hysteresis in the symbol grounding we get 61 undesired changes of the truth value. If we use a hysteresis with the size of σ then we reduce the number of changes to 17. An increase of the size of the hysteresis to $\Delta/2$ reduces the changes to zero. This clearly shows the benefit of the use of a hysteresis. But the size of the hysteresis is always a trade-off between stability and reactivity of the system. One has to take care that the size of the hysteresis does not exceed an adequate level. Otherwise, the system will lose its reactivity.

Figure 3 shows position measurements for an object within a dynamic experiment. The robot was placed 4000 mm away from the yellow goal. We recorded a series of position measurements while the robot performed a full rotation around its vertical axis. Positions are shown in the robots local coordinate system. The robot is located in the origin and the positive r-axis points to the front of the robot. If there were no inaccuracy in the vision system and the motion of the robot then the position measurements would lie on a perfect circle and the distance measurements to the object would remain constant. But the real measurements are affected by errors. There are three major reasons for these errors. First there is noise from the vision system. Furthermore, there is inaccuracy in the tracking of the object with the Kalman filter. This effect causes the tangential drift and the discontinuity of the measurements. Finally, the imperfect geometric calibration of the camera causes a deformation of the hypothetic circle. Distances to objects in the rear appear shorter in the camera as distances to objects in the front.

# Meas.	μ	σ	Δ	n			
				h=0	h=σ	h=$\frac{\Delta}{3}$	h=$\frac{\Delta}{2}$
	mm	mm	mm				
329	4789	485	1937	4	3	3	1

Table 2: Number of undesired change n of predicate *inReach* for the yellow goal for rotating distance measurements with different sizes h for the hysteresis.

Table 2 shows the evaluation of the above position measurements. The position measurements were converted to distances by calculating the Euclidean distance. Without using a hysteresis there are four undesired changes in the truth value of the *inReach* predicate for the yellow goal. The Table shows that with an increasing size of the hysteresis the number of undesired

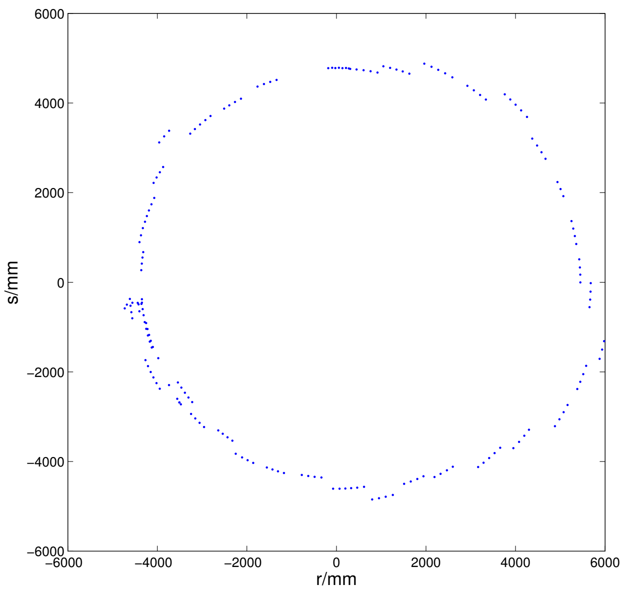

Figure 3: Position measurement for a static object 4000 mm away from the robot while the robot rotates. Positions are shown in the robots local coordinate system.

changes decreases to one. Please note that because all predicates are initialized with *false* there is always one change in the truth value even if the predicate is always correctly evaluated *true*.

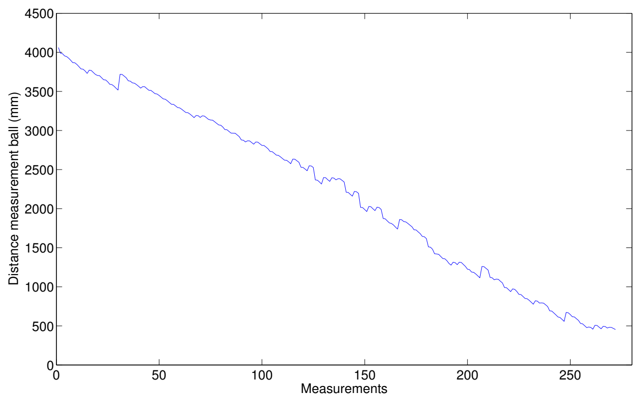

Figure 4: A sequence of consecutive distance measurement for a static object while the robot directly approaches the object.

Figure 4 shows the results of another dynamic experiment. In the experiment the robot was placed 4000 mm away from the ball and directly facing it. We recorded the distance measurements to the ball while the robot was directly approaching it. If we assume again a perfect perception then distances are supposed to monotonically decrease. The Figure clearly shows that this is not the case in our real experiment due to the imperfect perception. The evaluation of the symbol grounding without hysteresis for this experiment reports 3 undesired changes of the truth value of the predicate *inReach* for the ball. We calculated the mean and standard deviation of the differences of succeeding distance measurements, but we only considered cases in which the measured distance increased. The mean was 51 mm and the standard deviation was 8.8 mm. If we used a hysteresis with a size of σ, there were no undesired changes anymore.

The results of the above experiments in the real world show that quantitative perception is always affected by noise, changes in the environment and other inaccuracies. Therefore symbol grounding with simple thresholds can not be stable even the world does not change. This instability negatively affects the performance of the qualitative planning and reasoning process of an agent. Furthermore, the results show that the proposed symbol grounding with hysteresis is able to decrease the number of undesired changes of truth values of predicates to a minimum. This leads to an improvement of the stability of the robot's knowledge and increases the performance of the qualitative decision making process.

5 Open Issues

Although the use of a hysteresis in symbol grounding is justified by experimental results, there are still open questions concerning the proposed method.

There is no general answer to the question how the size of the hysteresis can be satisfactorily chosen. The size which is appropriate to sufficiently stabilize the symbol grounding while keeping the system reactive may differ from situation to situation. For example, the light conditions are always different and unpredictable. Furthermore, a more careful investigation should be done on the impact of the hysteresis on the reactivity of the system. A open question in this context is the definition of an appropriate evaluation criteria. A quick idea might be to play a dozed simulated games with and without the hysteresis and to compare the results, like goals scored or games won.

So far, we have not done any quantitative evaluation of how the symbol grounding with hysteresis influences the planning and plan execution in situations with slightly changes in the world. We assume that the hysteresis increases the performance of the plan execution as it can be compared to a commitment to follow a certain plan even if there are changes in the environment.

More research should be done on the conjunctions of predicates using hysteresis. Assume we use a conjunction of a large number of these predicates. If all measurements for predicates reach the upper boundary of their hysteresis the qualitative situation is the same as all measurements for predicates lie around the lower boundary. But the quantitative situations in the real world may substantially differ.

We use some predicates in different plans. Regardless of in which plan a predicate is used we use the same hysteresis size for the predicate. It might be desirable to use different hysteresis sizes for the same predicate in different situations in order to adjust the stability and reactivity of a predicate for a certain situation.

A small size for the hysteresis eliminates instabilities in the truth value without a significant decrease of the reactivity of the predicate. We need an even larger hysteresis if the inaccuracies in the perception become larger. But this fact negatively affects the reactivity of the system. It might be interesting if a smaller hysteresis is sufficient if more qualitative knowledge about how the world works is added to the reasoning.

6 Related Work and Conclusion

In [Reetz, 1999], an approach to action selection in Robotic Soccer is presented. The *action modules*, which are introduced in this work, have preconditions and invariants. The invariants can contain fewer conditions than the related preconditions in order to avoid oscillating behavior. The modules also have *activation factors* stating the utility of the action. These factors are situation independent, but are increased during the execution of an action module. This results in larger robustness of the behavior. In [Müller, 2000] a similar approach is used for gaining robustness. [Sachenbacher and Struss, 2001] presents a framework for automated qualitative abstraction of quantitative models. However, a complete knowledge of the quantitative model is required, whereas in our case only quantitative observations (which are incomplete and uncertain) are mapped. There are approaches which avoid the addressed problems in symbol grounding by the usage of reasoning with uncertainties like fuzzy logic or probabilistic networks. However, these approaches require different models and modeling processes.

In this paper we addressed the problem of symbol grounding in applications with a very high degree of (mostly unpredictable) interactions. We introduced the concept of predicate hysteresis to overcome some of the corresponding problems that occur in practice. We further described empirical results we obtained when using predicate hysteresis for symbol grounding on our robots. The outcome of the predicate hysteresis substantially improved the behavior of the robots. We further discussed open issues and future research directions.

References

[Fikes and Nilsson, 1972] Richard. E. Fikes and Nils J. Nilsson. Strips: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2(3-4):189–208, 1972.

[Fraser *et al.*, 2004] Gordon Fraser, Gerald Steinbauer, and Franz Wotawa. Application of qualitative reasoning to robotic soccer. In *Working Papers of the 18th International Workshop on Qualitative Reasoning (QR-04)*, 2004.

[Fraser *et al.*, 2005] G. Fraser, G. Steinbauer, and F. Wotawa. Plan Execution in Dynamic Environments. In *Proc. of the 18th Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems, IEA/AIE*, 2005.

[Fraser, 2003] G. Fraser. AI-Planning System for Robotic Soccer. Master's thesis, Graz University of Technology, 2003.

[Maybeck, 1990] P. S. Maybeck. The Kalman filter, An introduction to concepts. In *Autonomous Robot Vehicles*, pages 194–204, 1990.

[Müller, 2000] K. Müller. Roboterfußball: Multiagentensystem. Master's thesis, Universität Freiburg, 2000. In German.

[Reetz, 1999] C. Reetz. Aktionsauswahl in dynamischen Umgebungen am Beispiel Roboterfußball. Master's thesis, Universität Freiburg, 1999. In German.

[Sachenbacher and Struss, 2001] M. Sachenbacher and P. Struss. AQUA: A Framework for Automated Qualitative Abstraction. In *Working Papers of the 15th International Workshop on Qualitative Reasoning (QR-01)*, 2001.