

Qualitative World Models for Soccer Robots

Stefan Schiffer, Alexander Ferrein, and Gerhard Lakemeyer

Knowledge-Based Systems Group
Computer Science Department
RWTH Aachen University
Aachen, Germany
{schiffer,ferrein,gerhard}@cs.rwth-aachen.de

Abstract. Until now world models in robotic soccer have been mainly quantitative in nature, consisting of fine-grained (numerical) estimates of player positions, ball trajectories, and the like. In contrast, the concepts used in human soccer are largely qualitative. Moving to qualitative world models also for robots has the advantage that it drastically reduces the space of possible game situations that need to be considered and, provided the concepts correspond to those in human soccer theory, it eases the task of agent specification for the designer. In this paper we propose qualitative representations using ideas from spatial cognition and employing Voronoi diagrams. We also discuss how reasoning with these representations is achieved within our underlying agent programming framework.

1 Introduction

Until now world models in robotic soccer have been mainly quantitative in nature, consisting of fine-grained (numerical) estimates of player positions, ball trajectories, and the like. In contrast, the concepts used in human soccer are largely qualitative. Moving to qualitative world models also for robots has the advantage that it drastically reduces the space of possible game situations that need to be considered. Provided the concepts correspond to those in (human) soccer theory, it also eases the task of agent specification for the designer.

For example, if we use abstract to positional information like *front-left*, many similar game situations are represented by the same qualitative values whereas all these situations would differ in terms of their numerical values. This is useful, for example, when we formulate the preconditions required to initiate a tactical move. With a qualitative description we cover multiple similar settings, making the specification applicable in many circumstances. We also ease the specification process since we are able to use terms that are much closer to the natural language descriptions commonly used in human soccer theory. Besides, in the majority of cases a tactical instruction just cannot be formulated with precise positions but instead always refers to a set of positions denoted by a qualitative abstraction of regions such as *front* or *left*.

Dylla et al. [9] were among the first to address the question of how insights from human soccer theory can be applied when specifying the behavior of soccer


robots. Their proposal is to analyze existing moves from human soccer theory as, for example, described in [17] and to adapt these moves to the abilities of the respective robotic soccer leagues. They identified requirements needed to adapt existing moves to a soccer robot team one of which, they state, is that the robots in the team have to build a qualitative world model. The reason for this is that human soccer knowledge, which is often represented in the form of diagrams, has an inherent qualitative nature. To encode the moves which are most often depicted only in a prototypical fashion a qualitative world representation is needed to formalize behaviors for cooperative team play.

There already exists work on using qualitative information for autonomous agents. For example, Stolzenburg et al. [21] compared methods how to intercept the ball in the soccer simulator; one of the methods used qualitative abstracted ball coordinates. In [11], Fraser et al. describe the *inReach* predicate for the ball distance while employing a hysteresis function, Stone et al. [22] apply Reinforcement Learning in the Simulation League using some handpicked qualitative predicates for state space abstraction. While in these approaches only some qualitative aspects focusing on a particular task are chosen, none of these fulfill all requirements that are necessary for a complete qualitative framework.


In this paper we follow the ideas by Dylla et al. [9]. Continuing the work on using human soccer knowledge for robotic soccer, we present qualitative enhancements to a world model particularly suited for robots in the MIDDLE SIZE LEAGUE, where up to five robots per team play on a $8m \times 12m$ indoor soccer field. However, the enhancements could be applied to any other ROBOCUPSoccer league just as well. Drawing on previous work in the area of spatial cognition, we present an approach to representing positions on the field qualitatively. We also discuss models for further information needed to transfer human knowledge on soccer such as to decide when a pass can be played to a team-mate. This is done in the context of the logic-based action language READYLOG, a variant of GOLOG [16], which we use for an exemplary specification of the soccer move “kick-off”.

When using a qualitatively abstracted world model, an important issue is drawing appropriate inferences. A simple example is to derive the result of *distance_near + distance_far*. One option would be to apply one of the existing qualitative spatial calculi (cf. [4] for a survey of existing calculi). However, this currently does not seem feasible not only for computational reasons. Instead, we propose a hybrid quantitative-qualitative representation of the respective world model information, which allows for a limited form of reasoning about qualitative world model predicates, that seems expressive enough for most soccer applications. To be able to do so we have methods for re-quantifying the qualitative values to their numerical counterparts.

In the next Section we define our qualitative world model. In Section 3 we briefly introduce the language READYLOG and we show how it can be used to specify abstract soccer moves. In Section 4, we present an example illustrating our approach to reasoning about the predicates in our qualitatively enhanced world model. Section 5 concludes the paper.



(a) The combination of the distance and the orientation relation



(b) A point p defined in polar coordinates

Fig. 1. The combination of distance and orientation relations compared to the polar coordinate system.

2 Qualitative Representations

In the following we present the models which we use to abstract the quantitative data gathered from the sensors of the robots and stored in the quantitative world model to a qualitative representation of the world.

2.1 Positional Information

In [3], Clementini, Felici, and Hernandez present a unified framework which allows for qualitative representation of positional information. The framework combines an orientation and a distance relation. The position of a *primary object* is represented by a pair of distance and orientation relations with respect to a *reference object*. Both relations depend on a so-called *frame of reference* which accounts for several factors like the size of objects and different points of view.

The framework also features basic reasoning capabilities such as the composition of spatial relations as well as switching between different frames of reference. Unfortunately, the reasoning features provided are not guaranteed to yield unique¹ results, for instance, for the composition of qualitative terms. However, as for the ROBOCUP soccer domain we depend on unambiguous outcomes of such compositions since they are needed to instruct the robot.

From a quantitative point of view, the combined description of a position with this model can be seen as the representation of a point in polar coordinates. A point p in polar coordinates is defined by the distance r from the origin to this point and the angle φ measured from the horizontal x -axis to the line from the origin to p in the counter-clockwise direction. Thus, the position of a point p is described as (r, φ) . This description directly corresponds to the combination of the distance and the orientation relation. We illustrate this in Figure 1. This correspondence is of particular interest concerning our hybrid approach to reasoning which we will discuss in detail in Section 4.

The number of subdivisions, that is the level of granularity within the qualitative description of both distance and orientation can be chosen freely. In this

¹ By *unique* we mean results which contain exactly one relation.




Fig. 2. Semantic regions on the playing field. Figure (a) shows the orientation grid taken from [12]. Figure (b) shows the grid embedded into a soccer field. The resulting semantic regions on the playing field are shown in Figure (c).

paper we restrict ourselves to one level with eight distinctions although it is possible and might be of benefit to have multiple levels. For a recent approach to qualitative orientation with adjustable granularity see [19].

Beside the relative positional information described so far positions are also used in a global frame of reference. A qualitative concept which is applied frequently is that of semantic regions on the playing field. These regions are often used as tactical positions corresponding to player roles. To model semantic regions in the sense of global positioning we employ a well-known approach to qualitative representation of positional information proposed by Freksa and Zimmermann in [12, 13]. Qualitative orientation information in two-dimensional space is given by the relation between a vector and a point. The vector consists of a start point A and an end point B . It represents the orientation of a possible movement. Now, imagine a line through A and B and two further lines, one orthogonally going through A and B each. These three lines form an *orientation grid* which has the form of a *double-cross*. Different positions of an additional third point C can then be described with respect to this grid. Altogether the grid leads to 15 different orientation relations.

The grid used for the representation of orientational alignment is based on a (movement) vector. Unfortunately, we do not have an explicit movement in the context of global positioning on the playing field. On the other hand, we can regard the direction of play as a vector. From a tactical point of view one of the main objectives in a game is to advance from a defensive situation in the team's own half to an offensive one in the opponent's half. Thus, we can take the center of each team's half as the start and end point of an imaginary vector. If we place this vector onto the playing field we can relate semantic positions on the field to the orientation relations provided by Freksa's grid. This yields a subdivision of the field into 15 regions. The regions and their derivation are depicted in Figure 2.

Dividing the grid by the field's horizontal and vertical axes results in *zones* for the length of the field and in *sides* for the width. This system roughly corresponds to the Cartesian coordinate system; zones and sides form perpendicular axes with the zones corresponding to the x -axis and the sides corresponding to the y -axis.

Thus, we can still specify an object’s position in a coordinate system like manner, but by using zones and sides we achieve a qualitative abstraction. The analogy to the Cartesian coordinate system is of importance to our hybrid approach to reasoning again.

In the following we are going to elaborate on possible applications of the qualitative approaches to positional information we just presented.

2.2 Derived Predicates

Besides the greater comfort within the specification of tactical patterns we can use the zone and side information to build further qualitative predicates. Consider, for instance, the course of a soccer game. The overall positioning on the pitch is, besides the ball possession, one of the fundamental indicators upon which to classify whether a team is currently in a defensive or an offensive situation. Moreover, it is an important information whether the game’s focus is on one of the sides of the playing field or in the center.

With the zone information we can provide a predicate which we call *game setting*. It expresses where on the pitch the main activity of play is. It can be used to derive positioning instructions or simply to call appropriate sub-procedures in an agent’s program. We compute the game setting by a weighted sum of the zone indices of all players and of the ball: $w_{team} \cdot \sum_{i \in team} zone(i) + w_{opp} \cdot \sum_{j \in opps} zone(j) + w_{ball} \cdot zone(ball)$. The center of the possible values lies at a value of *zero* which states that the focus of play is located in the midfield. To be able to perform a classification of the game’s focal point in terms of a situation being *offensive*, *balanced*, or *defensive* we need to establish a threshold. Upon this threshold we can decide to which of the above class the current situation belongs. Feasible values for this threshold as well as for the weights for each of the three different objects within the above formula were found empirically in real world experiments. They can, however, also be learned which could lead to more adequate results. Analogously to the game setting we can determine the gist of play with respect to the sides of the pitch. We call this the *game edge*. The game edge renders useful, for instance, if we need to decide which side of the pitch is less occupied and can thus be used to advance into the opponent’s field half with a lower risk of being attacked.

2.3 Reachability

Apart from the static semantic regions there are further aspects which can be useful to determine. In particular, we are interested in dynamic spatial properties such as *reachability* of different kinds, which is central for the description and the execution of tactical patterns in soccer. For a discussion of the different forms of reachability we refer to [9]. For all these different reachability relations the individual abilities of a single robot or agent are relevant since even within the same league players of different teams may have unequal capabilities in terms of speed and mobility. Nevertheless, there are some general properties which hold for all players despite their physical abilities.

We consider the concept of Voronoi diagrams to be applicable for modeling a simplified version of the reachability relations required. A Voronoi diagram $V(S)$ of a set S of n point sites is the partitioning of a plane with n points into n convex polygons such that each polygon contains exactly one point and every point in the given polygon is closer to its central point than any other. For a more detailed account on Voronoi diagrams and their dual, the Delaunay triangulation $DT(S)$, see e.g. [1].

We make use of Voronoi diagrams and their dual the Delaunay triangulation to model reachability as they separate the field into non-intersecting regions and we get a connection graph between the players. We take the players as point sites in the plane and construct $V(S)$ and $DT(S)$ with the Euclidean distance thereupon. Then, the Voronoi region of each player is the set of points closer to this player than to any other player. Furthermore, in the dual of the Voronoi diagram $V(S)$, the Delaunay triangulation $DT(S)$, two players' point sites are connected if and only if they share a common boundary in the Voronoi diagram. Our idea of modeling reachability with the aid of $V(S)$ and $DT(S)$ is to consider players to be reachable to each other if their Voronoi regions share a common boundary, that is, they are connected in the Delaunay triangulation $DT(S)$. Fig. 3 depicts the Delaunay triangulation and the Voronoi regions for the geometric structure of several players on a soccer playing field. The yellow lines represent the triangulation, the green and red shaded regions correspond to the Voronoi regions of the attacking team and the defending team, respectively.




Fig. 3: Voronoi diagram and Delaunay triangulation of a soccer situation

2.4 Free Space

The notion of *free space* is another important aspect which can frequently be found in the description of spatial settings and tactical patterns in soccer. The term *free space* denotes an area which is not occupied by any of the players of the opposing team. Consider the Voronoi diagram being constructed upon all opponent players on the field. That is, each opponent player corresponds to the center of a Voronoi region. The Voronoi edges are formed by points that are equally far away from the two players the edge is between. The Voronoi vertices reflect points in the plane which have the maximal possible distance to even three or more of the surrounding players. These vertices directly match our interpretation of positions in free space as we sketched it above. Please note that we include the four corner points of the playing field into the Voronoi diagram. We need at least three Voronoi regions to obtain a Voronoi vertex. Including the corner points yields a minimum of four regions thus guaranteeing that we do not have less than one Voronoi vertex.

We provide two distinct ways to employ these vertices. First, we consider a classification request. Given a point on the playing field, we can ask how 'free'

this point is. To answer such a request we compute the point’s distance to the nearest point site in the opponent’s Voronoi diagram as well as its distance to the nearest Voronoi vertex. The ratio between these two values is a good criterion on how ‘free’ the given point is because it reflects whether the point is closer to a free position or to an opponent player. Since the ratio does not reflect the absolute distance to an opponent we additionally take a minimal distance into account which has to be exceeded. Second, we can answer inquiries for free point positions. Most times it is reasonable to specify a region of interest in which to search for a free position. For simplicity we assume that we specify a region of interest by a position in the pitch’s coordinate system (along with a maximal distance to search up to). Given this query, we determine the nearest Voronoi vertex to the position. If the distance between the query point and the vertex is lower than the maximal distance we return the vertex’ position. Otherwise, we return the position of a point lying on a line starting at the query point going into the direction of the nearest Voronoi vertex.

2.5 Additional Concepts

Within the course of a soccer game it is a vital piece of information whether or not one’s team is in possession of the ball. That is, again, an information we can provide by utilizing the structure of Voronoi diagrams. The simplest way to answer the question of ball possession is to check if the ball is located in the Voronoi region of a player who belongs to one’s own team. This is, of course, not always correct. For example, if the player whose Voronoi cell the ball belongs to is not facing the ball it might be the case that another player who has a greater distance to the ball but who is directly facing it can reach it more quickly. It is, however, possible to take this additional information into account and to refine the predicate accordingly.

As an additional qualitative predicate of particular interest in the soccer context we now consider something we call *passway vacancy*. We denote a qualitatively abstracted classification of the amount of space available along a potential pass way by this predicate. That is to say, we classify the degree of exposure of a line segment going from point P_{start} to point P_{end} by examining possible points of interception. We derive our classification by considering a ratio on how likely an interception is. Consider a straight line from P_{start} to P_{end} . We compute the minimal distance of each opposing player to this line, that is either the length of a line perpendicular to the pass way or the distance to the pass way’s nearest end point. Further, we compute the distance from each opponent to the starting point of the pass way. Then, we calculate the ratio between these two values. That is to say, we determine if the opponent is so close to the pass way that it can intercept a ball passed along the pass way.

3 Using Readylog For Behavior Specifications

READYLOG [10], a variant of GOLOG, is based on Reiter’s Situation Calculus [18,20], a second-order language for reasoning about actions and their effects.

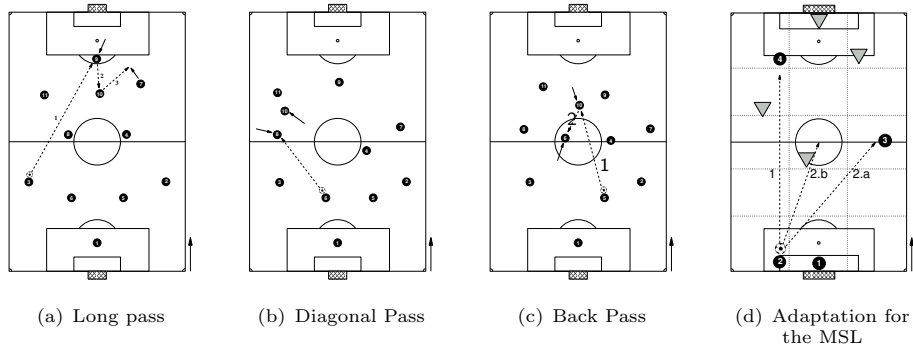


Fig. 4. Example for the “build-up play” move.

Changes in the world are only due to actions so that a situation is completely described by the history of actions starting in some initial situation. Properties of the world are described by *fluents*, which are situation-dependent predicates and functions. For each fluent the user defines a successor state axiom specifying precisely which value the fluent takes on after performing an action. These, together with precondition axioms for each action, axioms for the initial situation, foundational and unique names axioms, form a so-called *basic action theory* [20].

GOLOG has emerged as an expressive language in recent years. It has imperative control constructs such as loops (*while*), conditionals [16] (*if...then*), and (recursive) procedures (*proc(name(parameters), body)*), but also less standard constructs like the nondeterministic choice of actions (\mid). Extensions exist for dealing with continuous change [15] and concurrency [5], allowing for exogenous and sensing actions [6] and probabilistic projections into the future [14]. Another extension provides the facility to do decision-theoretic planning [2] which involves Markov Decision Processes (MDPs) (*solve(p, h)*, where p is a GOLOG program, h is the MDP’s solution horizon). READYLOG integrates these and more features in one agent programming framework [10].

To encode the behavior one has to specify a domain axiomatization including the actions the robot can perform together with their effects, and the fluents which describe the properties of the world like the ball position. Examples of domain descriptions for the soccer domain can be found in [7, 8].

3.1 A Soccer Move Example

We now specify the soccer move “kick-off” in READYLOG and we show that our qualitative world model (READYWORLD) supports the specification. We adapted three possible ways to build up a play as discussed in [17].

The first way to build up play is with a long pass (Fig. 4(a)). We immediately notice that the term *long* is one of the coarse, qualitative notions we need to establish in order to adapt human soccer theory for our autonomous soccer agents. We could also formulate this as passing the ball from a *back* position to a *front* position on the playing field. The second way to build up play is

```

proc build_up_play_defender,
  if haveBall(ownNumber) then
    getFreeSide(offense,FreeSide); getPassPartner(offense,FreeSide,PassPartner);
    solve(
      if  $\neg$ isKickable(ownNumber) then
        interceptBall
      else if isPassReachable(ownNumber,PassPartner) then
        passTo(ownNumber,PassPartner)
      endif
    )
    endif
  | pickBest(bestSide,{leftSide, middleSide, rightSide}
    dribbleTo(ownNumber,middleZone,bestSide)
    | kickTo(ownNumber,middleZone,bestSide) ) /* end of pickBest */
  | interceptBall; kickTo( ownNumber, middleZone, middleSide ),
  3, func.Reward ) /* end solve with horizon 3 */
  else
    interceptBall
  endif
endproc

```

Fig. 5. The build-up play program for the defender.

with a diagonal pass as depicted in Fig. 4(b). This time, the term *diagonal* is of qualitative nature. Diagonal means passing to the side being opposite to the current one. Fig. 4(c) shows the last possibility to build up play which is with a *deep* pass (dashed line labeled with 1) followed by a subsequent *back* pass (dashed line labeled with 2). The term *deep* is used to denote the space *behind* or *in between* a group of opponent players. The endpoint of such a pass has to be the most *free* position available *in between* or *behind* the group of opponents.

We now try to adapt as much of these descriptions as possible by integrating their most essential parts into one pattern. All three possibilities have in common that the ball is located in the *back* part on the pitch. According to a role ontology it is a player currently having a defensive role which is about to initiate the pattern to build up play. We already characterized the possibility of a long pass as 'bringing' the ball to the *front* part of the pitch. Therefore, in this case the agent chooses to pass to a teammate who is located in the attacking zone. For the two other possibilities the pass' destination is the midfield. The agent can either make a diagonal pass, that is the case if the target position is on the *opposite* side of the field, or it can simply pass to a free area on the same side or in the pitch's center. To illustrate our adaptation of the build up play patterns for the MIDDLE SIZE LEAGUE we depicted a diagram similar to the ones in [17] in Figure 4(d).

Figure 5 shows a program in our action language READYLOG capturing the above example. Note that this specification contains several qualitative elements such as *middleZone*, *leftSide*, and *offense* as well as qualitative predicates such as *isPassReachable*. With our qualitative world model we are able to simply transfer the qualitative notions from the specification in [17]. Moreover, the use of qualitative terms and predicates makes the program applicable in many game situations.

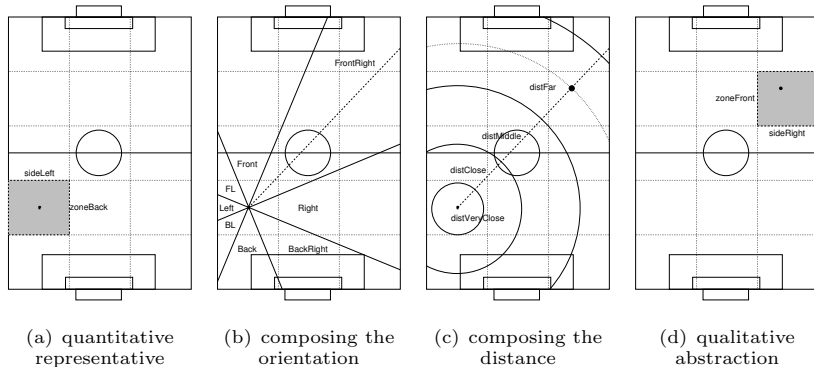


Fig. 6. An exemplary projection of qualitative values in READYLOG.

4 Reasoning

When a robot executes a READYLOG program like the one in Figure 5, it needs to evaluate different courses of action and choose the most appropriate one. This happens every time a *solve*-construct is encountered. Roughly, the robot evaluates the different alternatives (nondeterministic actions separated by '*'*' and nondeterministic choice of arguments (*pickBest*)) and chooses the one that maximizes expected utility in a decision-theoretic fashion (see [2, 10] for details). The evaluation of one alternative involves projecting the effects of the actions it contains into the future, starting from the current (qualitative) world model.

For the purposes of this paper, the main question is how moving-actions involving qualitative spatial terms are projected. For example, suppose the robot is currently in region $[zoneBack, sideLeft]$ (Figure 6(a)). What should the effect of the action $go(front-right, far)$ be? Depending on the exact position of the robot, it could end up in any of 9 zones on the field.² While this may still be manageable, things quickly get out of hand when we want to project a sequence of actions. In the worst case, the eventual outcome would span all 15 zones, that is, we lose all information about where the robot might end up. Moreover, projecting all possibilities is too costly, as the robot needs to make a decision on what to do next in the order of less than a second.

To avoid these complications and to remain computationally efficient, we simplify the problem in the following drastic way: to project the path of the robot, we simply translate the qualitative information back to numeric (geometric) values, taking as representative the mid-points of zones, sides, directions, and distances. The composition can be computed straight-forwardly using Euclidean geometry. The numerical end-result is then converted back to a qualitative description in the same way we perform the qualitative abstraction in the first place. Figure 6 illustrates this for the action $go(front-right, far)$ starting in region $[zoneBack, sideLeft]$. There are several advantages to this. Rather than having to entertain all possibilities, we only need to compute one. Using mid-

² See also [3] for a discussion of ambiguities in their approach to qualitative representations of positional information.

points also reduces the error in a reasonable way for practical purposes. Perhaps most importantly, we can use the existing projection mechanism of READYLOG, which requires that the effects of atomic actions are deterministic (in the case of a goto-action, the effect must be a unique location).³ The method is clearly sound, as the result is among those which a purely qualitative reasoner would obtain. It is not complete, as there may be cases where the computed path would take the robot outside the field and hence render the action illegal, even though other legal solutions may exist. Being incomplete is not a big problem in this application. It is more important to obtain a reasonable approximation fast in most circumstances. Also, plans often do not survive very long in soccer, that is, they are often aborted because the world has changed in a way that makes the current plan invalid. Hence it is not worth spending too much effort on figuring out what to do. Finally, we remark that, once a course of action has been chosen, we employ the same method that we use during reasoning to compute actual robot trajectories, using the real robot position as the starting location.

5 Conclusions

In this paper we proposed a qualitative spatial world model for soccer playing robots, combining earlier work on semantic regions with that on orientation and distance relations. In addition, we used Voronoi diagrams to provide us with a notion of reachability, which is important in the soccer domain. Computing robot trajectories from ambiguous qualitative descriptions was achieved by mapping the qualitative terms to unique geometric representatives.

The current reasoning scheme is admittedly somewhat ad hoc and was chosen mainly for efficiency reasons. One refinement would be to consider more than one trajectory. Also a comparison with existing spatial calculi [4] is needed. On the practical side, while the above ideas are fully implemented, we need to carry out more experiments to see how the qualitative approach fares in real games.

Acknowledgment

This work was supported by the German National Science Foundation (DFG) in the Priority Program 1125, *Cooperating Teams of Mobile Robots in Dynamic Environments* and a grant by the NRW Ministry of Education and Research (MSWF).

References

1. F. Aurenhammer and R. Klein. Voronoi diagrams. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*. Elsevier Science Publishers, 2000.
2. C. Boutilier, R. Reiter, M. Soutchanski, and S. Thrun. Decision-theoretic, high-level agent progr. in the situation calculus. In *Proc. AAAI'00*. AAAI Press, 2000.
3. E. Clementini, P. D. Felice, and D. Hernandez. Qualitative representation of positional information. *Artificial Intelligence*, 95(2):317–356, 1997.

³ Note that this is different from nondeterministic actions, which are complex, that is, made up of a number of primitive actions.

4. A. G. Cohn and S. M. Hazarika. Qualitative Spatial Representation and Reasoning: An Overview. *Fundamenta Informaticae*, 46(1-2):1–29, 2001.
5. G. De Giacomo, Y. Lésperance, and H. J. Levesque. ConGolog, A concurrent programming language based on situation calculus. *Artificial Intelligence*, 121(1–2):109–169, 2000.
6. G. De Giacomo and H. Levesque. An incremental interpreter for high-level programs with sensing. In H. J. Levesque and F. Pirri, editors, *Logical Foundation for Cognitive Agents: Contributions in Honor of Ray Reiter*, pages 86–102. Springer, Berlin, 1999.
7. F. Dylla, A. Ferrein, and G. Lakemeyer. Acting and deliberating using golog in robotic soccer – a hybrid architecture. In *Proc. CogRob02*. AAAI Press, 2002.
8. F. Dylla, A. Ferrein, and G. Lakemeyer. Specifying multirobot coordination in ICPGolog – from simulation towards real robots. In *AOS-4 at IJCAI-03*, 2003.
9. F. Dylla, A. Ferrein, G. Lakemeyer, J. Murray, O. Obst, T. Röfer, F. Stolzenburg, U. Visser, and T. Wagner. Towards a League-Independent Qualitative Soccer Theory for RoboCup. In *RoboCup 2004: Robot World Cup VIII*. Springer, 2005.
10. A. Ferrein, C. Fritz, and G. Lakemeyer. On-line Decision-Theoretic Golog for Unpredictable Domains. In *Proc. of 27th German Conference on Artificial Intelligence*, pages 322–336. Springer, 2004.
11. G. Fraser, G. Steinbauer, and F. Wotawa. Application of Qualitative Reasoning to Robotic Soccer. In *18th Int. Workshop on Qualitative Reasoning*, 2004.
12. C. Freksa. Using Orientation Information for Qualitative Spatial Reasoning. In A. U. Frank, I. Campari, and U. Formentini, editors, *Theories and methods of spatio-temporal reasoning in geographic space*. Springer, Berlin, 1992.
13. C. Freksa and K. Zimmermann. On the utilization of spatial structures for cognitively plausible and efficient reasoning. In *IEEE International Conference on Systems Man and Cybernetics*, pages 261–266, Chicago, 1992.
14. H. Grosskreutz. Probabilistic projection and belief update in the pgolog framework. In *Proc. 2nd Int. Cognitive Robotics Workshop*, pages pages 34–41. 2000.
15. H. Grosskreutz and G. Lakemeyer. cc-Golog – An Action Language with Continuous Change. *Logic Journal of the IGPL*, 2002.
16. H. J. Levesque, R. Reiter, Y. Lesperance, F. Lin, and R. B. Scherl. GOLOG: A logic programming language for dynamic domains. *Journal of Logic Programming*, 31(1-3):59–83, 1997.
17. M. Lucchesi. *Coaching the 3-4-1-2 and 4-2-3-1*. Reedswain Publishing, 2001.
18. J. McCarthy. Situations, Actions and Causal Laws. Technical report, Stanford University, 1963.
19. R. Moratz, F. Dylla, and L. Frommberger. A Relative Orientation Algebra with Adjustable Granularity. In *Proceedings of the Workshop on Agents in Real-Time and Dynamic Environments (IJCAI 05)*, 2005.
20. R. Reiter. On knowledge-based programming with sensing in the situation calculus. *ACM Transactions on Computational Logic (TOCL)*, 2(4):433–457, October 2001.
21. F. Stolzenburg, O. Obst, and J. Murray. Qualitative Velocity and Ball Interception. In M. Jarke, J. Köhler, and G. Lakemeyer, editors, *Proc. KI-2002*, LNAI, pages 283–298, Berlin, Heidelberg, New York, 2002. Springer.
22. P. Stone, P. Riley, and M. Veloso. The CMUnited-99 champion simulator team. In Veloso, Pagello, and Kitano, editors, *RoboCup-99: Robot Soccer World Cup III*. Springer, 2000.