Vision-Servoed Localization and Behavior-Based Planning for an Autonomous Quadruped Legged Robot

Manuela Veloso, Elly Winner, Scott Lenser, James Bruce, and Tucker Balch
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213-3891
mmv@cs.cmu.edu

Abstract

Planning actions for real robots in dynamic and uncertain environments is a challenging problem. Using a complete model of the world is not viable and an integration of deliberation and behavior-based reactive planning is most appropriate for goal achievement and uncertainty handling. This paper reports on our successful development of a system integrating perception, planning, and action for the Sony quadruped legged robots. We consider the robotic soccer task, as Sony provided the robots to us specifically for the RoboCup robotic soccer competitions. The quadruped legged robots are fully autonomous and thus must have onboard vision, localization and agent behavior. We briefly present our perception algorithm that does automated color classification and tracks colored blobs in real time. We then briefly introduce our Sensor Resetting Localization (SRL) algorithm which is an extension of Monte Carlo Localization. Vision and localization provide the state input for action selection. In addressing this planning challenge, we created a robust and sensible behavior scheme for the robot that effectively handles dynamic changes in the accuracy of the perceived information. We developed a utility-based system for using localization information. Finally, we have devised several special built-in plans to deal with times when urgent action is needed and the robot cannot afford collecting accurate perception information. We present results using the real robots demonstrating the success of the algorithms. Our team of Sony quadruped legged robots, CMTrio-99, won all but one of its games in RoboCup-99, and was awarded third place in the competition.

Introduction

The robots used in this research were generously provided by Sony (Fujita *et al.* 1999) to be applied to the specific domain of robotic soccer. The robots are the same as the commercial AIBO robots, but they are made available to us with slightly different hardware and programmable. The robot consists of a quadraped designed to look like a small dog. The robot is approximately 30cm long and 30cm tall including the head. The neck and four legs each have 3 degrees of freedom. The neck can pan almost 90° to each side, allowing the robot to scan around the field for markers. Figure 1 shows a picture of the dog pushing a ball.

Figure 1: The Sony quadruped robot dog with a ball.

All teams in the RoboCup-99 legged robot league used this same hardware platform. The robots are autonomous, and have onboard cameras. The onboard processor provides image processing, localization and control. The robots are not remotely controlled in any way, and as of now, no communication is possible in this multi-robot system. The only state information available for decision making comes from the robot's onboard colored vision camera and from sensors which report on the state of the robot's body.

The soccer game consists of two ten-minute halves, each begun with a kickoff. At each kickoff, the ball is placed in the center of the field, and each team may position its robots on its own side of the field. Each team consists of three robots. Like our team last year, CMTrio-98 (Veloso & Uther 1999), and most of the other eight RoboCup-99 teams, we addressed the multi-robot aspect of this domain by assigning different behaviors to the robots, namely two attackers and one goaltender. No communication is available and the robots can only see each other through the color of their uniforms. No robot identity can be extracted. As

of now, our robot behaviors capture the team aspect of the domain through the different roles.

The acting world for these robots is a playing field of 280cm in length and 180cm in width. The goals are centered on either end of the field, and are each 60cm wide and 30cm tall. Six unique colored landmarks are placed around the edges of the field (one at each corner, and one on each side of the halfway line) to help the robots localize themselves on the field. Figure 2 shows a sketch of the playing field.

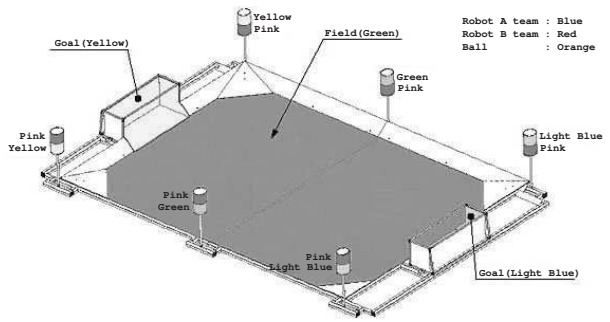

Figure 2: The playing field for the RoboCup-99 Sony legged robot league.

In this work, we address the challenges of building complete autonomous robots that can perform active perception and sensor-based planning. The robots perceive the world through vision, make decisions to achieve goals, and act by moving in the world.

We report in this paper on the three main components of our system integrating sensing, perception processing, and action selection, namely localization, vision, and behavior-based planning. We provide results within the particular RoboCup-99 domain and application.[1]

The vision algorithm is of crucial importance as it provides the perception information as the observable state. Our vision system robustly computes the distance and angle of the robot to the objects and assigns confidence values to its state identifications.

The preconditions of several behaviors require the knowledge of the position of the robot on the field. The localization algorithm is responsible for processing the visual information of the fixed colored landmarks of the field and outputting an (x, y) location of the robot.

Finally, our behavior-based planning approach interestingly provides the robot the ability to control its knowledge of the world. Behaviors range from being based almost solely on the visual information to depending on accurate localization information.

[1]Our extensive videos provide additional invaluable illustrative support to the contributions of this paper.

Vision

The vision system processes images captured by the robot's camera to report the locations of various objects of interest relative to the robot's current location. These include the ball, 6 unique location markers, two goals, teammates, and opponents. The features of the approach, as presented below, are:

1. **Image capture/classification:** images are captured in YUV color space, and each pixel is classified in hardware by predetermined color thresholds for up to 8 colors.

2. **Region segmenting:** pixels of each color are grouped together into connected regions.

3. **Region merging:** colored regions are merged together based on satisfaction of a minimum density for the merged region set for each color.

4. **Object filtering:** false positives are filtered out via specific geometric filters, and a confidence value is calculated for each object.

5. **Distance and transformation:** the angle and distance to detected objects are calculated relative to the image plane, and then mapped into ego-centric coordinates relative to the robot.

The onboard camera provides 88x60 images in the YUV space at about 15Hz. These are passed through a hardware color classifier to perform color classification in real-time based on learned thresholds.

When captured by the camera, each pixel's color is described as a 3-tuple of 8 bit values in YUV space. The color classifier then determines which color classes the pixel is a member of, based on a rectangular threshold for each class in the two chrominance dimensions (U,V). These thresholds can be set independently for every 8 values of intensity (Y). An example of the results of this classification is provided in Figure 3.

The final result of the color classification is a new image indicating color class membership rather than the raw captured camera colors. The 88x60 image has bits set for which classes, if any, a particular pixel is a member of. This is the input for the next step in the system, in which the connected regions of a particular color are determined.

The next stage is to compute a run length encoded (RLE) version for the classified image. This compression results in a substantial decrease in storage and processing requirements for subsequent steps. This is because the processing routines can operate on entire runs at a time, rather than individual pixels.

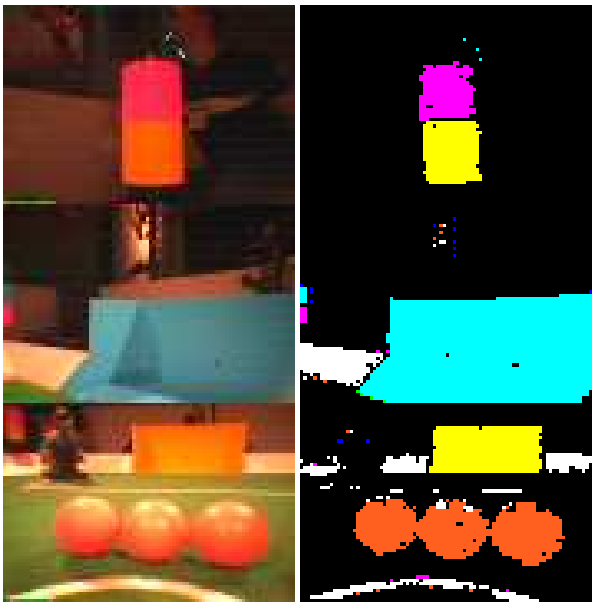

Figure 3: An example of our image classification on the right. The image on the left is a composite of objects: a position marker (top), a goal area (middle) and three soccer balls (bottom).

The region merging method employs a tree-based *union find* with path compression. This offers performance that is not only good in practice but also provides a hard algorithmic bound that is for all practical purposes linear. Initially, each run labels itself as its parent, resulting in a completely disjoint forest. The merging procedure produces a disjoint forest and a final pass compresses all of the paths in the forest so that each run's parent field is pointing directly to the global parent. Now each set of runs pointing to a single parent uniquely identifying a connected region, so the parent field can be thought of as a label which is unique to each region (Bruce, Balch, & Veloso).

We next extract region information from the merged RLE map. The bounding box, centroid, and size of each region is calculated incrementally in a single pass over the forest data structure. This process could easily be extended to extract additional statistics, such as a convex hull or edge points list. The information currently extracted provides enough information for the higher level manipulations.

After the statistics have been calculated, the regions are separated by color into separate threaded linked lists in the region table. Finally, they are sorted by size so that later processing steps can deal with the larger (and presumably more important) blobs, and ignore relatively smaller ones which are most often the result of noise.

The next step attempts to deal with one of the shortcomings of object detection via connected color regions. Due to partial occlusion, specular highlights, or shadows, it is often the case that a single object is broken into a few separate but nearby regions. A single row of pixels not in the same color class as the rest of the object is enough to break connectivity, even though the object may occupy many rows. In order to correct for cases when nearby regions are not connected but should be considered so, a density based merging scheme was employed. Density is represented here as the ratio of the number of pixels of the color class in the connected region to the overall area of its bounding box. By this measurement heuristic, two regions that have a small separation relative to their sizes will likely be merged, since they would tend to have relatively high density.

The next step is to finally calculate the location of the various objects given the colored regions. Various top down and geometric object filters are applied in each case to limit the occurrence of false positives, as well as serving the basis for confidence values.

For the ball, it is determined as the largest orange blob below the horizon. The confidence value is calculated as the error ratio of the density of the detected region and the actual density of a perfect circle. The distance is estimated as the distance required for a circular object to occupy the same area as the observed region. The field markers are detected as pink regions with green, cyan, or yellow regions nearby. The confidence is set as the error ratio of the difference between the squared distance between the centers of the regions and the area of each region (since they are ideally adjacent square patches, these two should be equal).

The colored blob on the marker indicates position along the length of the field. The relative elevation of the pink and colored regions disambiguates which side of the field the marker is on given the assumption that the robot's head is not upside-down. Thus the marker represented by a pair of regions can be uniquely determined. In case of multiple pairs which are determined to be the same marker, the one of maximal confidence is chosen. The distance to the marker is estimated from the distance between the centers of the two regions, since they are of known size.

The goals are detected similarly. They are the largest yellow or cyan regions with centers below the horizon. The distance measure is a very coarse approximation based on the angular height of the goal in the camera image, and the merging density is set to a very low value since many occlusions are possible for this large, low lying object. The confidence is estimated based on the difference in comparing the rel-

ative width and height in the image to the known ratio of the actual dimensions.

The final objects detected are opponents and teammates. Due to the multiple complicated markers present on each robot, no distance or confidence was estimated, but regions were presented in raw form as a list of patches. These simply indicate the possible presence of an opponent or teammate.

The final step in the vision system that needs to be mentioned is the transformation from image coordinates to an ego-centric coordinates. The vision system was found to perform well in practice, with a good detection rate and reasonably robust tolerance of the unmodeled noise experienced in a competition due to competitors and crowds. The distance metrics and confidence values also proved to be an advantage for localization and rational behavior in a highly noisy environment.

Localization

Our localization algorithm is based upon a classical Bayesian approach which updates the location of the robot in two stages, one for incorporating robot movements and one for incorporating sensor readings. This approach represents the location of the robot as a probability density over possible positions of the robot. In the CMTrio-98 localization algorithm, the probability density is represented using a grid based division of the pose space (Veloso & Uther 1999). Our localization algorithm, called Sensor Resetting Localization (SRL) (Lenser & Veloso), is based upon a popular approach called Monte Carlo Localization (MCL) which represents the probability density using a sampling approach.

Monte Carlo Localization(MCL) (Fox *et al.* 1999; Dellaert *et al.* 1999) represents the probability density for the location of the robot as a set of discrete samples. The density of samples within an area is proportional to the probability that the robot is in that area. Since the points are not distributed evenly across the entire locale space, MCL focusses computational resources where they are most needed to increase the resolution near the believed location of the robot. The position of the robot is calculated from these samples by taking their mean or some variant of mode. The uncertainty can be estimated by calculating the standard deviation of the samples. We encountered some problems implementing MCL for the robot dogs. MCL took too many samples to do global localization. With the number of samples we could actually run on the hardware, the samples were too spread out to localize the robot correctly. MCL also has problems dealing with large modelling errors.

SRL is motivated by the desire to use fewer samples, handle larger errors in modelling, and handle unmodelled movements. SRL adds a new step to the sensor update phase of the MCL algorithm. If the probability of the locale designated by the samples we have is low given the sensor readings $P(L|s)$, we replace some samples with samples drawn from the probability density given by the sensors $P(l|s)$. We call this sensor based resampling. The logic behind this step is that the average probability of a locale sample is approximately proportional to the probability that the locale sample set covers the actual location of the robot, i.e. the probability that we are where we think we are. This also mean that when tracking is working well, SRL behaves exactly the same as MCL.

This resetting step allows SRL to adapt to large systematic errors in movement by occasionally resetting itself. SRL is also able to recover from large unmodelled movements easily by using this same resetting methodology. Unexpected movements happen frequently in the robotic soccer domain we are working in due to collisions with the walls and other robots. Collisions are difficult to detect on our robots and thus cannot be modelled. We also incur teleportation due to application of the rules by the referee.

Movement update.
$P(l^{j+1}|m, l^j) = P(l^j)$ convolved $P(l'|m, l)$
[This stage is the same as Monte Carlo Localization]
1. foreach sample s in $P(l^j)$
(a) draw sample s' from $P(l'|m, s)$
(b) replace s with s'

Sensor update.
$P(l^{j+1}|s, l^j) = P(l^j) * P(l|s)/\alpha$ where α is a constant.
[Steps 1-5 of this stage are the same as MCL]
1. [optional step] replace some samples from $P(l^j)$ with random samples
2. foreach sample s in $P(l^j)$
(a) set weight of sample equal to probability of sensor reading, $w = P(l|s)$
3. foreach sample s in $P(l^j)$
(a) calculate and store the cumulative weight of all samples below the current sample (cw(s))
4. calculate total weight of all samples (tw)
5. foreach sample s' desired in $P(l^{j+1})$
(a) generate a random number(r) between 0 and tw
(b) using a binary search, find the sample with maximum cw(s) $<$ r
(c) add the sample found to $P(l^{j+1})$
6. calculate number of new samples, ns = (1 − avg_sample_prob/prob_threshold) * num_samples

7. if(ns $>$ 0) repeat ns times

(a) draw sample(s') from $P(l|s)$

(b) replace sample from $P(l^{j+1})$ with s'

Localization Capabilities

We tested SRL on the real robots using the parameters we used at RoboCup '99. We used 400 samples for all tests. In order to execute in real time, we were forced to ignore about 50% of the sensor readings. Due to inevitable changes in conditions between measuring model parameters and using them, the parameter for distance moved was off $\approx 25\%$, for angle of movement $\approx 10°$, and for amount of rotation $\approx .6°$/step. The deviations reported to the localization were 10% for movement and 15% for vision. We had the test robot run through a set trajectory of 156 steps while slowly turning it neck from side to side. We ran 5 times after 7 different numbers of steps had been executed. The final position of the robot was measured by hand for each run. We calculated the error in the mean position over time and the deviation the localization reported over time. We also calculated an interval in each dimension by taking the mean reported by the localization and adding/subtracting 2 standard deviations as reported by the localization. We then calculated the distance from this interval in each dimension which we refer to as interval error. We report both average interval error and root mean squared interval error. We feel that root mean squared interval is a more appropriate measure since it weights larger, more misleading errors more heavily. We also calculated the percentage of time that the actual location of the robot fell within the 3D box defined by the x,y, and θ intervals.

The table below shows the localization is accurate within about 10cm in x and y and 15° in θ despite the erroneous parameter values. The actual location of the robot is within the box most of the time and when it is outside the box, it is close to the box. The fact that the localization seldom gives misleading information is very important for making effective behaviors. The error in position and the deviation reported quickly converges to a steady level. The deviation tends to go up at the same time the error goes up which keeps the interval error low and avoids misleading output. In competition, we observed that the localization algorithm quickly resets itself when unmodelled errors such as being picked up occur. The actual performance of the localization during play tends to be worse than during testing since the robot spends much less time looking at the markers.

	x (mm)	y (mm)	theta (°)
average error	99.94	95.14	14.29
avg. interval error	15.18	4.91	2.07
rms interval error	34.92	13.94	3.82
in box percentage	74.29%	80.00%	57.14%

Behavior-Based Planning

The behavior of the robot is an especially difficult problem in this domain, in which the robot acts under uncertainty and must be able to quickly and gracefully improve and degrade its performance as the availability of localization information changes.

Because the localization system is reliant on visual identification of landmarks, in order to keep its localization information up-to-date, the robot must scan for landmarks. As the robot walks, the camera experiences pitch and roll, which causes the images it collects to change significantly from one frame to the next. Because of this, the vision system's estimate of angles degrades. Walking also creates lag between the time at which the vision frame was cpatured and the instant in the motion that the localization assumes the frame was captured at. Combined with processing time constraints, this makes the localization system produces more accurate estimates when the robot is stopped, so our algorithm has the robot stop moving while looking for landmarks. The process of stopping and scanning for landmarks usually takes the robot between 15 and 20 seconds.

Our approach provides the robot with the ability to control its knowledge of the world: in order to learn more about where it is, it can spend more time looking for landmarks. Although having more information helps the robot tremendously, soccer, like other dynamic domains, is time-critical, so every moment spent looking around is lost time. Opponents can use the robot's hesitation to their advantage.

We assumed that the vision system is correct and reliable in designing our behavior algorithms. Our strategy includes: i) a two-constraint system for utility-based thresholded localization, ii) an architecture that allows the robot to upgrade and degrade its performance quickly and gracefully, iii) behaviors that are reasonable even when the robot does not know where it is, and iv) several special localization-dependent behaviors which dramatically increase the robots' efficiency.

Control over State Knowledge

The robot must balance time spent localizing with time spent acting. One possible localization strategy, used by this year's team from LRP University in France (Bouchefra *et al.* 1999), involves localizing the

robot very infrequently, if at all. However, the benefits of accurate localization are significant.

We present a scheme that balances the time required to get accurate localization information with the benefits this information provides.

Utility-Based Thresholded Localization We found that if we defined what a "useful" amount of localization information was and allowed the robot to stop and look around every time its localization information fell below that level, the robot spent most of its time looking for landmarks. However, if we changed the definition of "useful" information to the point that the robot was able to act for a reasonable amount of time before stopping to look for landmarks, then its localization information was almost never good enough to use; in our tests, the robot frequently scored own-goals, or tried to push the ball into the walls of the field.

We use a system of two constraints to force the robot to act for long enough to avoid disrupting its behavior while also requiring that its localization information is accurate enough to use.

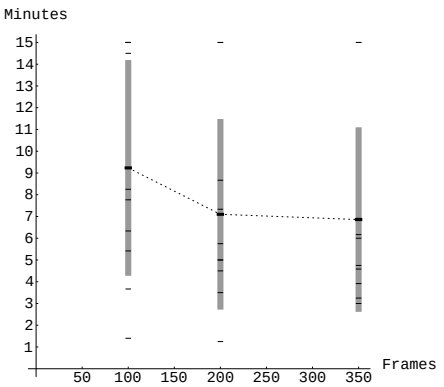

Figure 4: Time taken to score a goal versus how long we require the robot to act before looking for landmarks.

Constraint 1—Enforcing Action: Our first constraint is that the robot spend a certain amount of time acting before it stops to look for landmarks. This ensures that the robot does not spend all of its time localizing but still ha the opportuinity to benefit from localization.

Our scheme uses a counter to require the robot to act for a specified amount of time before looking for landmarks. The amount of time the robot must act before looking could depend on the confidence the robot has in its current localization information and on its current goals. In our scheme, however, it is invariant.

We require that the robot act for the time it takes the image module to process 350 frames of data, or about 40 seconds. Recall that stopping to look for landmarks takes the robot between 15 and 20 seconds, not counting the time it takes it to recover the ball afterwards. So we demand that it spend about 2/3 of its time acting. The results of our experiments, shown in Figure 4 show that this value is good (Winner & Veloso).

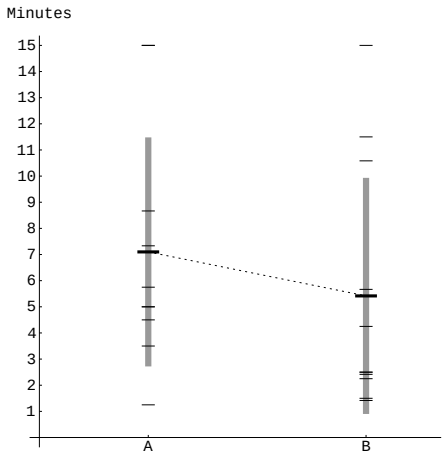

Figure 5: Time taken to score a goal versus how low we require the standard deviations of the localization values for θ (the robot's orientation) and x and y (its location on the field) to be before using them.

Constraint 2—Limited Localization: The second constraint is how accurate we demand the localization information to be. We measure accuracy with the standard deviations returned by the localization module. If the information is accurate enough, the robot should not stop to look for landmarks. But if it is not accurate enough, the robot should not use it. It is not immediately obvious how to balance allowing the robot to use localization information of questionable accuracy with the cost of having the robot ignore accurate localization information.

Our results, displayed in Figure 5 show that a "good" localization estimate of θ, or the angle of the robot on the field, have a standard deviation of 30° or less. A "good" localization estimate of x and y, the coordinates of the robot's location on the field, must have a standard deviation of less than 600mm, or just under one quarter of the field length.

Behaviors with No Need for Localization One of the best ways of finding a balance between localization information and time spent to acquire it is simply to avoid localizing when it is not necessary. We have

identified two times when it is unnecessary: when the robot has recently lost sight of the ball, and when it has no information at all about the location of the ball.

Recovering a Recently Lost Ball: The robot often loses sight of the ball while it is trying to manipulate it. This is usually because it has walked past the ball. It must search for the ball, but because it has just lost sight of the ball, more information is available for it to use.

Instead of incorporating a full-scale world view into the robot architecture, we implemented a very simple but extremely effective algorithm. When the robot loses sight of the ball, it first walks backwards for a while. If it has walked past the ball, this usually allows it to spot it. However, sometimes walking backwards makes the robot turn its body away from the ball. So, like our team last year, CMTrio-98 (Veloso & Uther 1999), this year's robots turn in the direction in which the ball was last seen. After this, the robot considers the ball lost and begins a random search for it.

Random Search: When the robot does not know where the ball is, it must wander the field to search for it. One way of searching for the ball is to build an "oriented" search, in which the robot uses localization information to systematically search each area of the field. This relies on accurate localization information and on a complete search of the field, both of which take a lot of time. Instead, we use a very simple algorithm that is much faster. We wrote a random search algorithm that, until it sees the ball, alternates between walking forward a random amount and turning a random amount. The random search avoids the cost and necessity of localizing.

Acting with Little Information

Frequently during games, the standard deviations of the robot's localization information are so high that the information should not be used. As explained previously, the robot should not stop and look every time its localization information is inaccurate. Therefore, we must make sure that it can act reasonably even when its localization information is not good enough to use.

We wrote an algorithm which allows the robot to score goals without information about the robot's angle and x and y location on the field. The algorithm is as follows:

```
Until see the goal,
  walk sideways to the right around the ball;
When see the goal,
  line up with goal and ball;
Walk forward into the ball.
```

This allows the robot to consistently score goals with no information from the landmark-based localization module at all. However, this takes the robot much longer than it does when the robot has such information.

Upgrading Performance

Not only must the robot be able to progress towards its goals when it does not have useful localization information, it must be able to improve its performance as soon as it gets such information. We used several different strategies to make sure that the robot's performance would be able to improve quickly as well as degrade gracefully as the availability of good information changed.

Localization-Dependent Performance Enhancements We have developed three performance enhancements that rely on localization information and that are robust and reliable even with noisy information (Winner & Veloso). The first helps the robot decide in which direction to circle around the ball, or whether to circle at all. The second allows the robot to skew its approach to the ball so that it doesn't have to spend time circling. The third allows the robot to score goals even if it is unable to see the goal itself.

Mode-Based Architecture We built a control architecture based on basic modes of behavior we identified as important. To switch between modes, the robot uses knowledge about basic features of its state, such as: i) whether it is paused, ii) into which goal it is trying to push the ball, iii) whether it is in possession of the ball, i.e. close to the ball, and iv) whether it knows where the goal is. Localization enhancements improve swtiching between modes and performance within a mode. The modes we have defined for the attacker are:

1. `Head Searching:` the robot is searching for the ball by turning its head;
2. `Searching:` the robot is searching for the ball by turning its head and its body;
3. `Approaching:` the robot is approaching the ball;
4. `Circling:` the robot is circling the ball;
5. `Scoring:` the robot is pushing the ball towards the goal;

The algorithm we use to switch among these modes is approximately as follows:

```
If robot sees ball and is not close to it,
  mode = Approaching;
If robot does not see ball and did recently,
  mode = Head Searching;
If robot does not see ball and has not recently,
  mode = Searching;
```

```
If robot does not know where goal is and
    is close to ball
  mode = Circling;
If robot knows where goal is and is close to ball,
  mode = Scoring.
```

We are able to optimize the performance of the low-level implementation of the modes by using localization information. By separating the high-level behavior from the low-level implementation, we ensure that the robot's high-level behaviors do not change frequently as the available information changes.

Special Cases—Urgent Action

In some cases, the balance between localization and action does not apply because immediate action is needed and localization takes too long. We have found shortcuts and compromises that allow the robot to perform as well as possible in these special cases without localization information.

Approaching the Ball Possession of the ball is a critical part of a soccer game. The team that possesses the ball more has an incredible advantage over its opponents. Therefore, in our strategy, when the robot sees the ball, it rushes towards it, not waiting to localize.

This strategy has negatives, clearly. If the robot does not know where it is on the field, it will not know what to do with the ball when it gets to it. Nevertheless, it is better for a robot to look around when it is in possession of the ball than when it is farther from the ball. When the robot is standing near the ball, it is blocking one side of the ball from visibility and attacks.

Kickoff We were surprised to discover, in the initial games of RoboCup-99, how much of an advantage is gained by winning the kickoff. When the ball moves to one side of the field, it is very difficult for the robots to move it to the other side of the field. In the RoboCup-99 games, the team that won the kickoff usually scored a goal, simply because the ball never emerged from the side of the field to which it was initially pushed.

Because of this, we developed and extremely aggressve kickoff that uses the fact that the robots start each kickoff facing the ball and the opponent goal. We allowed the robots to run with the ball without localizing at all for almost half the length of the field. Even when the error-prone motion of the robots causes them to stray far off course, the ball is usually driven into the opponents half before our robots are forced to localize.

This was a very important advantage in almost all of our games, since we were able to win most of the kickoffs. In the one game we lost, we had accidentally turned this feature off during the first half of the game, and our robots were no longer able to win the majority of the kickoffs. When we turned it back on in the second half, our robots again dominated the kickoffs, and were therefore able to score a goal and prevent the other team's robots from scoring any.

Goalie Another time when swift action is crucial is when a robot is playing the position of goaltender. However, this position also requires very accurate localization, since it is necessary for the goalie to be in the correct position in front of the goal. We found that our localization information was not accurate or fast enough to localize the goalie in front of the goal because of its reliance on the markers around the field.

By avoiding the landmark-based localization module altogether, we were able to find a way for the goaltender not only to avoid looking frequently at landmarks, but also to position itself more accurately in front of the goal. Our final algorithm is as follows:

```
Starting Position: Centered in front of the goal,
  facing the other side of the field.
Scan the horizon for the ball;
If the ball is seen, run straight after it;
If lose sight of the ball for more than 2 frames,
  turn until own goal is seen;
If see own goal, run towards largest area of goal seen
  until it fills visual field;
If own goal fills visual field
  turn until opposing goal is seen.
```

This final version of the goalie is extremely aggressive, and extremely successful. One of the main strengths of this algorithm is that it takes advantage of the special situation in which the goal tender finds itself—standing very close to one goal, and facing the other. Because the goals are the largest visual features on the field, it is easy to use them to localize this special position. Because the goalie pushes the ball far away from the goal, it usually has plenty of time to run back to the goal and turn around before the ball comes nearby.

Conclusion

In this paper, we reported on our work on controlling the soccer-playing Sony quadruped legged robots based on visual perception and probabilistic localization. We briefly described the vision and localization algorithms that allow for the state information to be gathered during execution of the game.

We then contributed a behavior-based planning approach that actively controls and balances the amount of localization information the robot has. The robot can score goals relying solely on the limited visual perception. The behaviors also employ as much of the localization information as is available and they upgrade and degrade performance gracefully as availability changes. In addition, the robots include deliber-

ative preset plans to deal with special cases in which urgent action is necessary and therefore cannot afford the time to gather accurate state information. We include results of tests that demonstrate the localization capabilities and support our parameter settings to control the amount of localization information.

Results from our matches in RoboCup-99 at IJCAI-99, Stockholm, also show our algorithms to be effective. Our team won all but one of its games, and the one it lost was lost by only one goal. Our team was the only one in this year's league to score goals against opposing teams and never to score a goal against itself. Our goaltender was the only one in this year's league to score a goal itself.

Acknowledgments: We would like to thank Sony for providing us with this remarkable platform for our research in the perception, planning, and action.

References

Bouchefra, K.; Hugel, V.; Blazevic, P.; Duhaut, D.; and Seghrouchni, A. 1999. Situated agents with reflexive behavior. In *Proceedings of IJCAI-99 Workshop on RoboCup*, 46–51.

Bruce, J.; Balch, T.; and Veloso, M. Fast color image segmentation using commodity hardware. *Submitted to ICRA-2000.*

Dellaert, F.; Fox, D.; Burgard, W.; and Thrun, S. 1999. Monte Carlo localization for mobile robots. In *Proceedings of IROS-99.*

Fox, D.; Burgard, W.; Dellaert, F.; and Thrun, S. 1999. Monte Carlo localization: Efficient position estimation for mobile robots. In *Proceedings of AAAI-99.*

Fujita, M.; Veloso, M.; Uther, W.; Asada, M.; Kitano, H.; Hugel, V.; Bonnin, P.; Bouramoue, J.-C.; and Blazevic, P. 1999. Vision, strategy, and localization using the Sony legged robots at RoboCup-98. *AI Magazine.*

Lenser, S., and Veloso, M. Sensor resetting localization for poorly modelled mobile robots. *ICRA-2000.*

Veloso, M., and Uther, W. 1999. The CMTrio-98 Sony legged robot team. In Asada, M., and Kitano, H., eds., *RoboCup-98: Robot Soccer World Cup II.* Berlin: Springer Verlag. 491–497.

Winner, E., and Veloso, M. Robust action under variable uncertainty: An algorithm for robotic soccer. *Submitted to Autonomous Agents-2000.*