

Learning and Using Models of Kicking Motions for Legged Robots

Sonia Chernova and Manuela Veloso
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
{soniac, mmv}@cs.cmu.edu

Abstract—Legged robots, such as the Sony AIBO, create opportunity to design rich motions to be executed in specific situations. In particular, teams involved in robot soccer RoboCup competitions have developed many different motions for kicking the ball. Designing effective motions and determining their effects is a challenging problem that is traditionally approached through a generate and test methodology. In this paper, we present a method we developed for learning the effects of kicking motions. Our procedure acquires models of the kicks in terms of key values that describe their effects on the ball’s trajectory, namely the angle and the distance reached. The successful automated acquisition of the models of different kicks is then followed by the incorporation of these models into the behaviors to select the most promising kick in a given state of the world. Using the robot soccer domain, we demonstrate that a robot that takes into account the learned predicted effects of its actions performs significantly better than its counterpart.

I. INTRODUCTION

Many different kicking motions for quadruped robots have been developed in recent years by the teams involved in the RoboCup competitions. These motions are designed to propel the ball in various directions with different speeds. As the number of available motions grows, the process of selecting which kick to use has become more complex.

Learning the effects of deterministic actions has been studied in classical planning (e.g. [1], [2]) where the learning algorithms extract the preconditions and effects of actions through experimentation under different world conditions. Reinforcement learning assumes that the environment is a Markov Decision Process and learns the model of the world, i.e., it learns the nondeterministic effects of the actions through experimentation [3], [4]. In this work, we learn the effects of the actions of our robot decoupled from complete task performance, as it is not feasible to assign reward directly to specific state and action pairs in a continuous execution sequence of the robot performing its usual robot soccer behavior.

We present a method for modeling the effects of the kicks in terms of several key values describing the ball’s trajectory. Specifically we analyze the angle of the ball’s trajectory, the distance traveled by the ball when actuated by the kick, and the success rate of the kick. We then incorporate these models into the behaviors to select the most promising kick in a given state of the world. Our results show that using this model the robot achieves its goals more effectively than a robot that does not take into account the predicted effects of its actions.

For data gathering we chose to use only the local sensors on the robot, mainly the color camera located in the head of the robot. As a result, these experiments can be run in any environment where the robot is able to localize itself without the need to setup any additional equipment. This method can be adapted to a variety of robot platforms where the task is to learn the effects of defined motions on objects in the environment.

We begin by providing background information and our motivation for pursuing this topic in Section II. The algorithms for modeling the angle of the ball’s trajectory and the strength of the kicks are discussed in Sections III and IV respectively. In Section V we discuss how these models can be incorporated into the behaviors to select the most effective kick. Experimental results comparing scoring performance with and without kick modeling are presented in Section VI, and our conclusions are presented in Section VII.

II. MOTIVATION

The robots used in this research are the Sony AIBO four-legged robots. Through several years working with these robots, we have developed a fully autonomous software system for soccer-playing robots. The work described in this paper focuses on how the robot can autonomously model the effects of its own motions, and use the derived model to select appropriate motions in the future.

The motions that we would like to model are the kicking motions that the robot uses to propel the ball while playing soccer. Our goal is to study the effects that each kick has on the location of the ball. In particular, we would like to represent the effect of the kick in terms of the expected displacement of the ball, and the angle of the ball’s trajectory.

Each of the robot’s kicks is encoded using frame-based motion, which describes the transitions of the body frame by frame by specifying a series of body, leg, and head positions and a time period for interpolating between one position and the next. Generally lasting only a few seconds, these motions are designed to be executed the same way every time. The Forward Arm and Hard Left Head Kick are shown in Figures 1 and 2 respectively.

Each robot is equipped with a color camera that is mounted into the head of the robot. The three degrees of freedom of the head, combined with an approximate 55° field of view of



Fig. 1.



Fig. 2.

the camera, allow the robot to track objects over a wide area in front of and next to the robot. The onboard camera will be the only sensor used in our analysis. It will be used to report the distance and angle of the ball relative to the robot, as well as locations of several known landmarks which will be used to triangulate the robot's position.

The accuracy of location estimates for various objects reported by the vision system varies with respect to distance and the movement rate of the camera. Since the camera is the only sensor used, we briefly discuss the accuracy of its measurements. Figure 3 compares the levels of noise in the sensor readings to five ball positions at two different camera movement rates. In Figure 3(a), the robot estimates the position of the ball while it is standing and the camera is still. In Figure 3(b) the robot reports estimates for the same ball locations while it is pacing in place causing the camera to move up and down. The results show that while the robot is stationary, the angle estimates to the ball are very reliable, with higher uncertainty in the distance estimate. Both distance and angle estimates become less reliable when the camera moves while the robot is pacing. The most accurate location estimates

are achieved when the robot is standing still a small distance away from the ball. A similar experiment using localization landmarks produced similar results.

III. TRAJECTORY ANGLE

The angle of the ball's trajectory relative to the direction the robot is facing is an important characteristic of all kicking motions. In this section we will describe an algorithm for estimating the angle of the trajectory for a variety of kicking motions using only the robot's camera.

In order to calculate the angle of the ball's trajectory we record the path of the ball over the period of 1 second (25 frames) immediately after the kick. There are two main benefits for analyzing this short segment of the trajectory. First, the ball has not yet moved far away from the robot and our estimates of the ball's position will be most accurate in this range. Second, the ball has the greatest velocity at this point and will travel the true path in which it was kicked. As the ball's velocity decreases, the ball tends to follow an unpredictable curve resulting from small imperfections in the ball's shape and irregularities of the surface. By studying the initial trajectory we avoid introducing this additional noise into the model.

By tracking the ball immediately after the kick, the robot is able to fit a regression line to the data and approximate the angle of the trajectory. Table-I shows the algorithm developed that allows the robot to perform this task autonomously.

The proposed algorithm can be executed in two modes, with and without human assistance for ball placement. As shown, the algorithm requires a human assistant to place the ball in front of the robot for each trial. This improves the consistency of the experiment by guaranteeing similar conditions for each trial. The same procedure can also be executed with the robot searching for and approaching the ball after each kick. Although completely autonomous, this method may not be as accurate if the robot is not able to approach the ball well in case of obstacles.

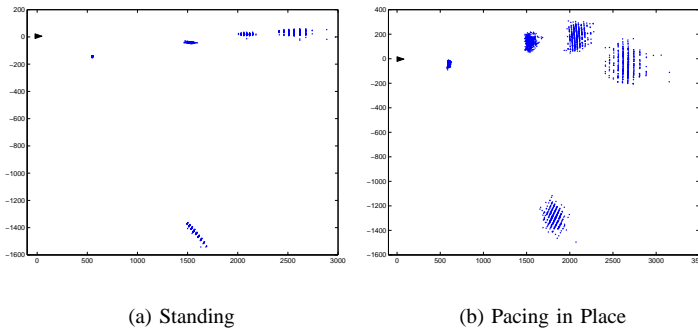


Fig. 3. Ball location estimates. Reported ball locations for five stationary balls at various distances and angles while the robot is standing or pacing in place. The location of the robot is marked by the black triangle.

Algorithm III.1: TRACKANGLE()

```

timeOfKick ← 0
while 1
  TRACKBALLWITHHEAD()
  if BallWithinKickingRange = true
    do
      then { KICK()
            timeOfKick ← currentTime
            if currentTime - timeOfKick > tdelay
              then { angle ← CALCANGFROMBALLLOCHIST()
                    output (angle)

```

TABLE I

COMPUTATION OF THE ANGLE OF BALL'S TRAJECTORY FROM AN INPUT OF THE ESTIMATED BALL DISTANCE AND ANGLE VALUES FROM VISION.

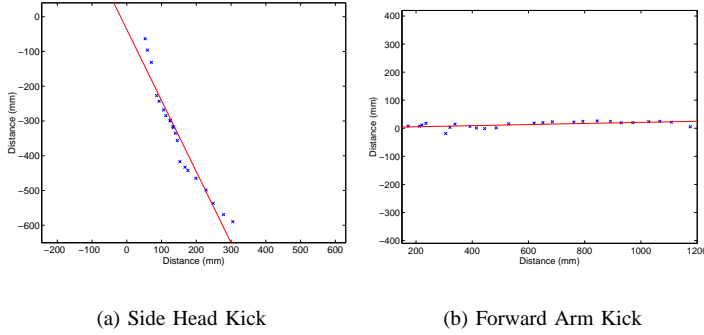


Fig. 4. Single trial analysis of two kicks. Each point represents the position of the ball relative to the robot in a single vision frame. A regression line is fitted to the points to estimate the angle of the ball's trajectory.

To assure that the robot was able to track the ball successfully, we require that at least 20 of the 25 polled frames contain information about the location of the ball. Figure 4 shows the angle analysis results of a single trial for the Forward Arm and side Head Kicks. Note that the regression line is much more sensitive to variations in the estimated angle measurement to the ball than to the estimated relative distance. Using the results from our analysis of reported ball locations while standing and pacing, we can conclude that the trajectory of the ball at such close range while the robot is not moving is approximated with very high accuracy.

In Figure III we summarize the results of angle analysis for the Forward Arm, Normal Left Head Kick and Normal Right Head Kick over 480 trials. The means of the the three kicks are 2.1°, 72.6°, and 55° respectively, with variances of 82.81°, 20.25°, and 31.36°.

IV. DISTANCE

The second attribute important in understanding the effects of the different kicking motions is the distance the ball travels, or the strength of the kick. In this section we will describe an algorithm for estimating the distance the ball travels, as well as calculating the average success rate of the kicking motion.

The robot is unable to track the entire trajectory of the ball because the ball travels beyond the robot's visual range for most of the kicks. Instead, our algorithm uses the final resting location of the ball relative to the original position of the robot




Fig. 5. Trajectory angle analysis results for 410 trials of the Left Head Kick, Forward Kick and Right Head Kick.

before the kick to estimate the strength of the kick. Table II shows the algorithm used to calculate the displacement of the ball after a kick.

The robot performs this analysis without any human assistance. Each trial takes approximately 1-2 minutes. Calculations of both the ball position relative to the robot, and the robot's own location relative to known landmarks are taken while the robot is standing in order to increase the accuracy of the measurements. When estimating the location of the ball the robot remains at a small distance in order to avoid accidentally bumping into and moving the ball.

In addition to estimating the strength of a particular kick, this algorithm can also be used to determine the success rate of the kicking motion. A kick is considered to have failed if proper contact is not made and the ball is moved only a few centimeters, if at all. Failed kicks can be detected easily using a simple distance threshold to distinguish between successful and unsuccessful trials. Detecting failed trials allows us to establish a reliability measure for each kick, as well as exclude these results from the analysis.

Figure 6 summarizes the results of distance analysis of the Normal and Hard Left Head Kicks. The hard head kick propels the ball much further, with some distances nearing 3.5 meters

Algorithm IV.1: TRACKDISTANCE()

```

while 1
  APPROACHBALL()
  KICKBALL()
  STANDANDLOCALIZE()
  initBallLoc ← currentRobotLoc
  FINDBALL()
  do
    APPROACHBALL()
    if ballDistance < 50cm
      STANDANDLOCALIZE()
      then { finBallLoc ← currentBallLoc
            ballDispVec ← finBallLoc - initBallLoc
            output (ballDispVec)

```

TABLE II

COMPUTATION OF A VECTOR REPRESENTING THE BALL'S DISPLACEMENT RELATIVE TO THE LOCATION OF THE KICK, GIVEN THE ESTIMATES OF THE BALL AND ROBOT LOCATIONS FROM VISION.

Kick	Angle Mean(deg)	Angle Variance(deg)	Dist Mean(m)	Dist Variance(m)	Success Rate
Forward	2.1	82.81	2.2	2.07	85%
Normal Head L.	72.6	20.25	1.48	0.33	98%
Normal Head R.	-70.4	31.36	1.48	0.33	98%
Hard Head L.	72.6	20.25	2.57	0.62	90%
Hard Head R.	-70.4	31.36	2.57	0.62	90%

TABLE III
THE LOOKUP TABLE.

with an average distance of 2.57 meters. The normal head kick has a range of at most 2 meters with an average of 1.48 meters.

The wide range of final locations for the ball shows the difficulty of modeling the effects of the kicks. In some trials the kick fails completely and the ball does not move at all, as can be seen for one of the trials of the Hard Head Kick where the ball's final position coincides with the location of the robot. In other trials the robot makes a strong contact with the ball but possibly with the wrong part of the body, or at the wrong angle, which results in an unpredicted trajectory for the ball. This can cause the ball to roll in the opposite direction than expected, or even to curve around behind the robot.

V. BEHAVIORS

We selected two specific attributes to model the effects of the kicking motions, the angle of the ball's trajectory and the distance traveled by the ball after the kick. We used the acquired data to build a model that represents each kick in terms of its effects on the ball. To incorporate the model into the behaviors we create a lookup table containing the attribute values for each kick. Table III is an example of such a table

containing five different kicks. Note that this table makes two small assumptions. Since the head kicking motions are symmetric in the left and right directions, we are making the assumption that the Left and Right Head Kicks have the same strength in both directions. The second assumption in the table, made because no angle data was gathered on the Hard Head Kick, is that the Hard and Normal Head Kicks have the same trajectory angle. Ideally both distance and angle values would be measured for every kick in the table.

The robot behaviors reference the lookup table to select the appropriate kick to use. When selecting a kick, the robot calculates the desired trajectory of the ball to the target goal, and uses a selection strategy to select the most appropriate kick. Different selection strategies can be developed for different situations by weighting the importance of some attributes over others. For example, if the robot is close to the goal, the angle of the ball's trajectory becomes more important than the strength of the kick, while from far away a stronger kick would be more desirable. Such preferences can easily be translated into numerical selection strategies and sets of rules for which strategy should be used.

Kicking motions can easily be added or removed from behaviors simply by editing the lookup table. If none of the kicks in the lookup table satisfy the current selection strategy, several behaviors can be sequenced together to achieve the desired effect. For example, the robot may chose to turn or dribble the ball to achieve a better scoring position.

VI. EXPERIMENTAL RESULTS

The presented kick selection algorithm was tested by comparing the performance of two robots running the code from CMPack'02, Carnegie Mellon's robot soccer team. On one robot the behavior system was modified to include the lookup table and selection algorithms described. The robots were tested on their ability to score a goal on an empty field without any opponents present. Testing in this manner guarantees that the data upon which the selection algorithm relies, mainly the location of the robot, is most accurate. Multiple robots would interfere with each other and push as they compete for the ball, which would effect the localization system. This would make it impossible to distinguish whether a poor kick was a result of poor kick selection, or simply because the robot was lost.

For each trial the robot begins at the goal line of its own goal, and the ball is placed at one of the four predefined points




Fig. 6. Distance analysis of the Normal and Hard Left Head Kicks. Each point represents the final resting position of the ball after a kick, relative to the initial position of the robot marked by the triangle.

that are unknown to the robot, see Figure 7.



Fig. 7. Experiment setup.

The robot’s performance is evaluated by recording the time it takes to score on the opponent goal. The four points chosen for the experiment are designed to test a variety of distances and angles to the target goal. For example Point1 is chosen to be far away but at a very direct angle to the goal, while Point4 is near the goal but at a very steep angle.

Each robot ran a total of 52 trials, 13 for each of the four points. Table IV summarizes the results of the experiment. For every point the robot using the presented selection algorithm scored faster, with an overall average improvement of 13 seconds. The statistical significance of the results was confirmed using the Wilcoxon Signed Rank test with a 0.05 significance level.

Point	CMPack’02	Modeling
Point1	56.7	39.8
Point2	42.5	27.2
Point3	76.5	60.0
Point4	55.0	52.0
Total	57.8	44.8

TABLE IV

PERFORMANCE COMPARISON OF CMPACK’02 VS THE PRESENTED KICK SELECTION ALGORITHM. VALUES REPRESENT MEAN TIME TO SCORE IN SECONDS, AVERAGED OVER 13 TRIALS PER POINT.

VII. CONCLUSION

We have presented a method for autonomously modeling the effects of kicking motions in terms of attributes describing the behavior of the ball. We then incorporated this model into the behaviors in the form of a lookup table or a motion library. This information was then used to select appropriate motions with various selection strategies. Using the robot soccer domain we have demonstrated that a robot which takes into account the predicted effects of its actions performs significantly better than its counterpart.

This algorithm extends to a wide range of tasks in which the robot must select the appropriate action to execute from a set of possible actions. Through observation of changes in

the state of the world, a model predicting the effects of each action can be learned, and used to make better informed action decisions in the future.

ACKNOWLEDGMENT

The authors wish to thanks Scott Lenser, Douglas Vail and James Bruce for their valuable contributions.

REFERENCES

- [1] Y. Gil, “Acquiring domain knowledge for planning by experimentation,” Ph.D. dissertation, School of Computer Science, Carnegie Mellon University, August 1992, available as technical report CMU-CS-92-175.
- [2] X. Wang, “Learning planning operators by observation and practice,” in *Proceedings of the Second International Conference on AI Planning Systems, AIPS-94*, Chicago, IL, June 1994, pp. 335–340.
- [3] L. P. Kaelbling, M. Littman, and A. Moore, “Reinforcement learning: A survey,” *Journal of Artificial Intelligence Research*, vol. 4, pp. 237–285, 1996.
- [4] M. Asada, S. Noda, S. Tawaratsumida, and K. Hosoda, “Purposive behavior acquisition for a real robot by vision-based reinforcement learning,” *Machine Learning*, vol. 23, no. 2-3, pp. 279–303, 1996. [Online]. Available: citeseer.nj.nec.com/article/asada94purposive.html