In Proceedings of RoboCup International Symposium, Singapore, June 2010.

Designing effective humanoid soccer goalies

Marcell Missura, Tobias Wilken, and Sven Behnke

University of Bonn,
Computer Science VI, Autonomous Intelligent Systems
Roemerstr. 164, 53117 Bonn, Germany
{missura,wilken,behnke}@cs.uni-bonn.de
http://ais.uni-bonn.de

Abstract. Most of the research related to the topic of falling strategies considers falling to be an unavoidable part of bipedal walking and is focused on developing strategies to avoid falls and to minimize mechanical damage. We take an alternative point of view and regard falling as a means to an end. We present our falling strategy for the specific case of a robot soccer goalie that deliberately jumps in front of a moving ball to prevent it from rolling into the goal. The jump decision is based on observed ball position, speed and direction of movement. We show how we implement a targeted falling into the appropriate direction, minimize the time from the jump decision to ground impact, and what solutions we developed to prevent mechanical damage. The presented falling technique was used in RoboCup Humanoid KidSize and TeenSize competitions and proved to be essential for winning.

Key words: humanoid robots, robot soccer, falling, motion generation, mechanical design

1 Introduction

Falling is an inevitable part of bipedal walking, especially in dynamic environments such as robot soccer games. However, falling can also be intentional. In highly dynamic sports players frequently decide to take risky actions that result in falling to the ground. Two soccer-related examples are kicking from an unstable position to attempt to move the ball towards the opponent goal even at the cost of falling, or jumping in front of the ball to prevent it from reaching the own goal. Unlike accidental falls, intentional falls do not strike the player by surprise. The location and time of the ground impact can be estimated more precisely and the player has more time to prepare a fall sequence to minimize the risk of damaging the body.

In this paper, we present our falling strategy for the specific case of a robot soccer goalie that deliberately jumps in front of a moving ball to prevent it from rolling into the goal. The jump decision is based on observed ball position, speed and direction of movement. We show how we designed a diving motion that does not damage the robot's body. The fall is accelerated to reach the ground as soon as possible and targeted to the left or right side depending on where we expect to catch the ball. The proposed method was used in the German Open 2009 and RoboCup 2009 competitions, where they proved to be essential for winning, for example the TeenSize Dribble & Kick competition in Graz 2009.

The remainder of the paper is organized as follows. In Section II we present a review of related work. In Section III our simple, trainable algorithm is presented that handles the jump decision. Section IV describes how we developed the diving motion for the goalie and explains the motion itself in detail. In Section V we show our mechanical solutions for damage avoidance and finally, in Section VI we present experimental results from RoboCup competitions.

2 Related Work

While the research of humanoid bipedal walking has been a hot topic for decades, the issue of falling has barely been addressed. Some work exists on the detection of situations that would lead to a fall [1-3] and on avoiding a fall altogether by stopping [4], squatting for a short while [3], or stepping into a capture region [5]. Results addressing actual falling strategies when the ground impact is inevitable, are scarce. The research groups of Fujiwara et al. [7-12] and Ogata et al. [13-14] have proposed UKEMI techniques that distribute the impact force over a sequence of designated impact points [10]. Forward and backward falling motions were optimized using inverted pendulum based models in [11-12]. Yun et al. [6] presented a different approach of changing the fall direction using targeted stepping and inertia shaping to avoid hitting an obstacle. In the closest related work, Ruiz del Solar et al. [15] also came to the conclusion of distributing the impact force onto multiple contact points and proposed a methodology to design fall sequences that can also be used for intentional falling. To the best of our knowledge, no research up to date is published on intentional falling, although it is applied by several teams in the KidSize league of the RoboCup humanoid soccer competitions.

3 The Jump Decision

The first thing the goalie does when it enters the field, or after getting up from the floor, is to use global localization on the soccer field to position itself on the center of the goal line looking into the direction of the opponent goal. When the position of the goalie is good enough, the robot stops, bends its knees to lower the CoM, and observes the field in this special goalie halt position that is lower than the halting posture of the field players. In order to protect the goal, the goalie tracks the ball and estimates the position and the velocity of the ball in egocentric coordinates that are used to calculate when and in which direction the goalie should jump. Currently, our robots are equiped with a 3 camera vision system. This is subject to change due to a rule update in the RoboCup Humanoid Soccer league that prohibits the use of more than two cameras. Our

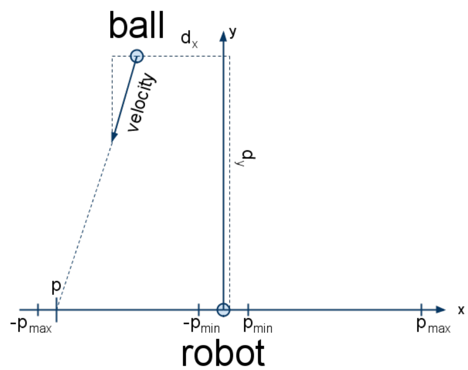

Fig. 1. A sketch of the egocentric view of the goalie and estimated contact point p of the ball with the lateral plane x. The goalie will only jump if the ball is expected to hit the goal between the goal posts (p_{max}) and with some distance to the feet of the goalie (p_{min}).

vision system processes approximately 25 frames per second distributed over 3 cameras. In the exceptional case of the goalie halt position, all but one cameras are switched off. This way the available 25 frames per second are concentrated on one camera only to boost the precision of the velocity estimations. The position of the ball is measured by inverting the ground projection onto the camera plane. The velocity of the ball is estimated by averaging the change of position in three consecutive frames.

Figure 1 illustrates the egocentric view of the goalie. We denote the lateral direction as x. Negative values are to the left and positive values to the right of the robot. The sagittal direction was labeled y with positive values to the front. The jump decision of the goalie is based on the estimated time and point of contact of the ball with the lateral plane.

The one-dimensional contact point p on the lateral plane is given by:

$$p = d_x - d_y \frac{v_x}{v_y}, \tag{1}$$

with d_x and d_y being the distance of the ball in lateral and sagittal direction relative to the goalie and (v_x, v_y) is the velocity of the ball. Only a certain range of contact points are interesting because they are located inside the goal area. Assuming the goalie is standing in the middle of the goal and $|p| > p_{max}$, the ball will entirely miss the goal on the left or the right side. On the other hand, if the contact point is too close to the goalie ($|p| < p_{min}$), the ball will most likely bounce off of the feet of the goalie. In this case instead of a diving motion the goalie squats down quickly holding its arms to the sides of its feet in order to achieve a larger blocking surface. We set p_{max} to half of the width of the goal and p_{min} to half of the width of the robot.

Apart from the point of contact, the time of contact is also an important determinant of the jump decision. If the ball is slow and takes long enough to reach the goal for the goalie to have enough time to position itself between the ball and the goal, there is no need to jump. The diving motion is associated with high costs, such as a certain time needed for the goalie to get back on its feet and to reposition itself in the goal. During this time the goalie is incapacitated and the goal is unprotected. In addition, the risk of mechanical damage is always present. We only want the goalie to perform the diving motion when it is necessary to block the ball. To determine the time of contact is a more difficult task than to calculate the contact point, because the ball does not travel with a constant velocity. In fact, the friction with the ground slows the ball down at a rate which is highly location dependent on the floor material and possibly uneven ground. Additionally, in the KidSize and the TeenSize leagues balls of different size and weight are used. Because an analytic calculation would be complex and possibly misleading, we developed a very simple trainable model. We collect training data by placing the goalie in the goal and do not allow it to move. Then we kick the ball towards the goal from various distances and random velocities. After each kick we indicate with a joystick wether the goalie should have jumped or not and label a whole series of observations. After we collected enough data, we reduce the state space to two dimensions by discarding the x components of the ball and velocity observations and use libsvm [16] to train a support vector machine to separate the two classes. Figure 2 shows an example. The blue colored region shows the cases where the goalie should not jump, the red region indicates the class where the goalie should jump. The training data is colored according to the manually assigned labels. It is noticable that red and blue observations overlap slighty. This is because for the border cases even for humans it is diffi-

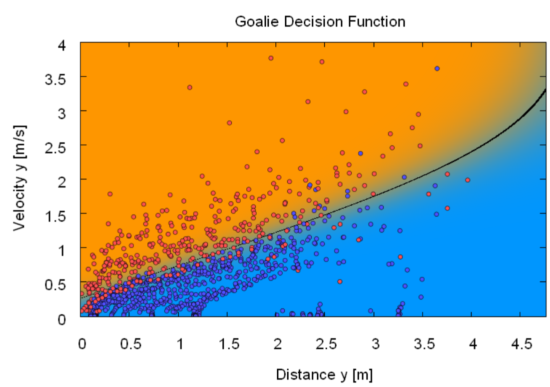

Fig. 2. Classification of ball observations into two categories: the goalie should jump (red) and the goalie should not jump (blue). The black line marks the border between the two classes, which is fine tunable with one parameter.

cult to decide if the goalie should have jumped or not. Furthermore it can be seen that starting from approximately two meters, the noise in the position and velocity estimations increases significantly with distance. The black line marks the border between the two classes. A nice property of libsvm is that it provides a probability estimate of the certainty that an observation belongs to a class. By default we seperate the classes where an observation belongs to each class with a probability of 0.5. Later on this threshold can be altered to manually fine tune the "sensitivity" of the goalie with a single paramater in case the goalie occasionally lets a ball pass or behaves too jumpy.

4 The Goalie Motion

The diving motion of the goalie is comparable to an inverted pendulum. If started in a perfectly upright position at 90°, neither the pendulum, nor the goalie will fall. The more the starting angle of the pendulum deviates from the vertical, the faster it reaches the ground.

Excluding any kind of active pulling towards the ground, we can regard a free falling point mass dropped from the height of the robot as the lower bound of possible falling times. The time the free falling point mass takes to reach the ground from a height of 0.6 m and an initial velocity of zero can be calculated using Newton's law.

$$h(t) = h_0 - \frac{1}{2}\, g\, t^2 \tag{2}$$

Setting $h(t) = 0\ m$, $h_0 = 0.6\ m$, $g = 9.81\ m/s^2$, and solving for t we obtain a falling time of approximately 0.35 seconds. We are using an inverted pendulum as a benchmark that approximates the size of the robot's body (0.6 m), starting from a small angle of 0.1 radians. We determined the falling time of the pendulum numerically with Euler's method using

$$\ddot{\phi}(t + \Delta t) = -\frac{g}{l}\, cos\, \phi(t) \tag{3}$$

$$h(t) = l\, sin\, \phi(t) \tag{4}$$

and setting $\dot{\phi}(0) = 0$, $\phi(0) = \pi/2 - 0.1$, $l = 0.6\ m$ and $g = 9.81\ m/s^2$. ϕ describes the angle of the pendulum with $\phi = \pi/2$ being the vertical and angles are increasing counter clockwise. We used a very small time step of $\Delta t = 10^{-6}\ s$. According to the simulation it takes 0.87 seconds for our benchmark pendulum to reach a height of zero. When inspecting the height of the falling pendulum, as shown in Figure 3, one can clearly see that the extra time is needed mainly for the first degrees of falling. Once the pendulum passes an angle of 45°, which is equivalent to a height of approximately 0.4 m, the remaining fall is almost as fast as free falling. Apart from avoiding damage, the challenge in designing a diving motion for the goalie is to accelerate the fall with an initialization motion, comparable to starting an inverted pendulum at angles deviating from the

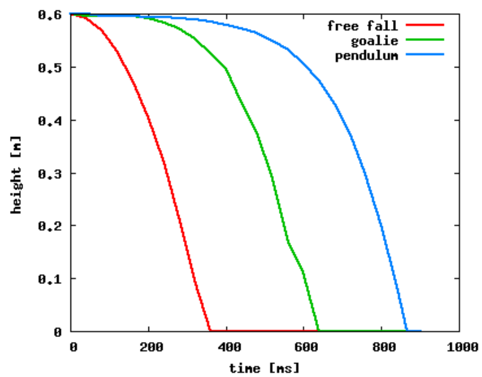

Fig. 3. Time series of the fall height of the goalie motion (middle) compared to a free falling body (left) and an inverted pendulum started at an angle of 0.1 rad (right).

vertical as much as possible.

The diving motion of our goalie is driven by a short sequence of key frames as depicted in Figure 4. Due to the symmetry of our robots, the same motion can be applied for jumps to the left or the right by mirroring the key frames respectively. First, in key frame a) the goalie stands motionless in the goal and observes the field in a special goalie halt position that is bent more in the knees than the posture of the field players. When it is time to jump, the motion starts in key frame b) with a sideways hip swing using the hip and the ankle joints to accelerate the torso towards the yielding leg. We named the leg towards which the robot is falling the yielding leg, the other one the support leg. A few moments later the yielding leg starts to shorten in key frame c), while the feet are constantly rolling to support the fall of the robot to the side. Before the support leg loses ground contact, it is extended and moved inwards closer to the yielding leg in key frame d). We found that the resulting rotation of the torso reduces the time until the ground impact, while the sudden leg extension increases the

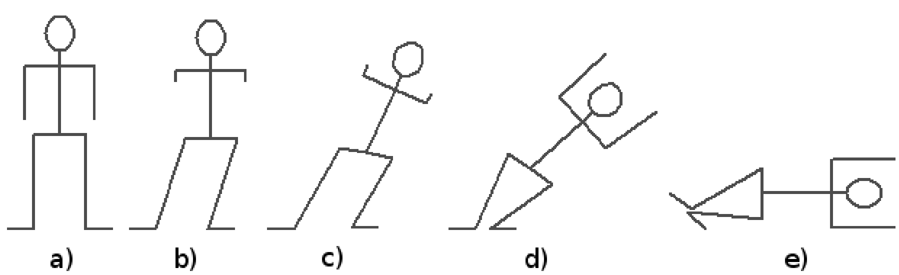

Fig. 4. The key frames of the goalie motion.

blocking distance. Simultaneously, the arms are lifted up high above the head with the maximum speed the servos allow. The arms reach their target position just in time before they touch the ground in key frame e), so that when the goalie lands, it can touch the goal posts with its arms. Finally, when lying on the floor, the goalie performs a "sweeping" motion with the yielding leg to clear a ball that may have been blocked and is now lying dangerously close to the goal. When performing the diving motion, the special goalie halt position plays a key role. As mentioned before, the posture of the goalie is lower than the default standing position of our field players for two reasons. First, the CoM is lower and it takes less time for it to reach the ground. And second, the bent knees of the robot give room for a longer, more powerful push with the support leg. Earlier versions of our goalie started in full upright position and were significantly slower. Furthermore, the arms of the goalie are already raised a bit when standing in the halt position. This is important for the arms to reach their target position above the head in time. The shoulder servos are not fast enough to complete the full 180 degrees from the bottom to the top in the short time available. If the arms are not up at the time of the ground impact, the goalie may be injured by the fall and the blocking distance would also be shorter since the arms would not reach as far when the robot is down on the floor.

Our diving motion takes 0.64 seconds from the jump decision to ground contact. We have determined this time by recording the goalie motion several times with a digital camera and counting the frames from when the goalie started to move until the robot obviously touched the ground. Multiplying the frame rate of the camera (12 fps) with the average number of frames, we obtain the total duration of the motion. Figure 3 shows the results as the height of a reference point in all three experiments depending on time. The reference point of the goalie is on the top of the head. The height of this reference point was determined by counting the pixels in the frames of a video from a base line up to the top of the head and converting them into centimeters. Please note that these numbers were obtained by experiments with our KidSize robot. The actual diving times of our TeenSize robot may differ slightly.

We developed the goalie motion in a simulation with a three dimensional physical model of our KidSize robot. Only after we achieved promising results in the simulation we made the first attempt on a real robot. The motion worked close to the expectations without damaging the hardware, so we continued to optimize the motion in a real environment. When the motion was first applied to a real robot after it was developed in simulation, it was significantly slower with 0.8 s until ground contact. By making small changes to the keyframes manually and testing it on the real robot we have been able to improve the timing. With the current duration of 0.64 s it is equivalent of starting an inverted pendulum of 0.6 m length at approximately 14 degrees (0.25 rad).

5 Preventing Mechanical Damage

Although the fall of the goalie is intentional, preventing mechanical damage is still a crucial aspect. The intention of the fall, in contrast to accidental falling, makes it easier to prepare a damage minimizing pose in time before the ground impact. Earlier versions of our goalie motion turned the robot's torso towards the ground and used the arms to soften the impact of the fall, as can be seen in Figure 5. Our KidSize robots now have less sensitive shoulder joints after a degree of freedom in lateral direction was removed. We found that raising the elastic arms above the head is sufficient to avoid damage to the shoulder joints and designed a new motion without the more time-consuming torso twist. The first point of impact occurs on the upper arm as can be seen in Figure 6. The arms are somewhat flexible and are padded with soft guards, which cushion the impact and also protect the head that contains the cameras. With only one degree of freedom in the shoulder in sagittal direction, the ground forces are orthogonal to the rotational plane of the servo and the gears are not stressed by the impact. Nevertheless, triggered by the last keyframe of the motion shortly before the impact, all joints of the robot are completely relaxed to effectively protect the gears of the servos, the most sensitive parts when it comes to falling. Except for the joints, the robot is well protected against falls, since the aluminum frame is sturdy enough to withstand the ground impact from a low height such as the size of the KidSize models.

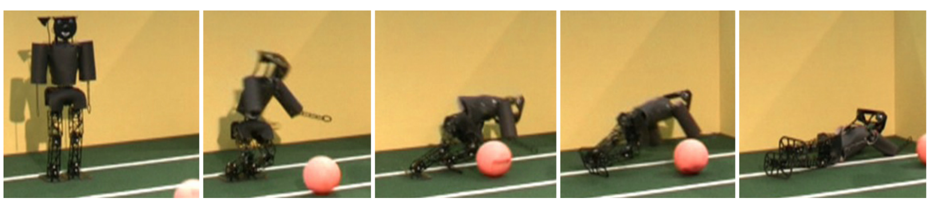

Fig. 5. TeenSize 2007 robot Bodo protects the goal by twisting its torso and leaning on its arms to land the fall.

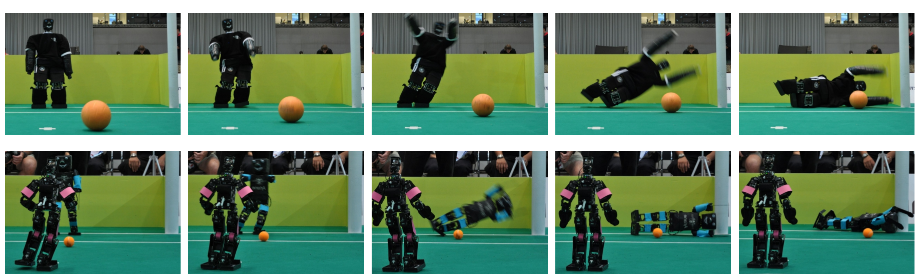

Fig. 6. TeenSize robot Dynaped (top row) and KidSize robot Ariane (bottom row) take a dive for the ball during RoboCup 2009 in Graz. The attacker in the bottom row was a penalty striker from team CIT Brains.

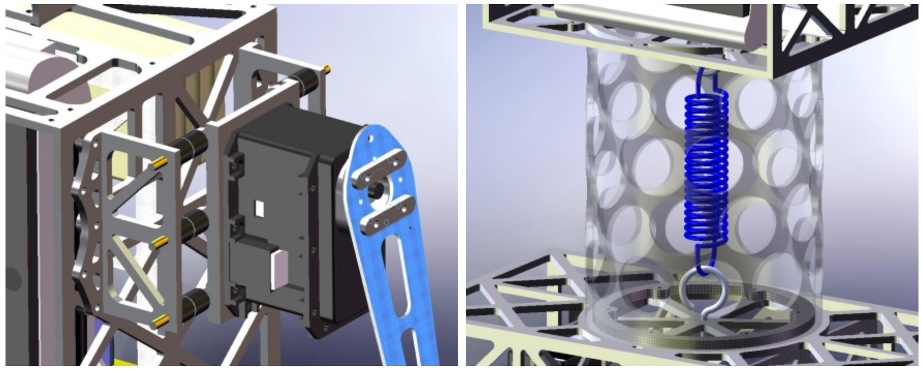

Fig. 7. TeenSize robot Dynaped's mechanical precautions against damage resulting from falls. The shoulder (left) is fixed to the frame by six flexible struts made from hard rubber. The hip (right) is a pull linkage with a spring that holds the torso in place.

As already mentioned, our TeenSize model Dynaped was performing the same goalie motion. Dynaped is significantly larger and heavier than our KidSize robots. In addition to relaxing the joints shortly before the impact, we also designed special mechanical precautions to avoid damage to the frame of the robot, as shown in Figure 7. Dynaped's elastic arms are completely padded with a thick layer of foam and have no joints apart from the saggital joint in the shoulder. The shoulder joint itself is attached to the torso with a flexible connection. Six struts made from hard rubber hold the shoulder in place and yield even to weak forces. A rigid connection would likely bend or break sooner or later when the goalie repeatedly hits the floor with its arm. Furthermore, the torso of the robot is attached to the hip with a pull linkage that is equiped with a strong spring. The right part of Figure 7 shows this hip construction. The spring loaded linkage works much like a mechanical fuse. The trunk cylinder shown semi transparently is not mechanically fastened to the hip. It is loosely positioned on a ring that prevents the cylinder from sliding sideways. The blue spring holds the cylinder firmly in place, but it yields to strong-enough forces. We are using a spring with a spring constant of 1.753 N/mm expanded by 36.6 mm in the base position. This results in a minimum force of 64 N exerted by the spring at all times that holds the torso firmly in place. Using the screws at both ends of the spring the initial tension can be changed easily. We found the appropriate setting by manually adjusting the screws until the spring force was high enough to keep the torso from shaking during walking.

Figure 8 sketches an analysis of the ground impact. The expected point of attack of the ground reaction force is the shoulder joint. The height of the torso from the base of the cylinder up to the middle of the shoulder joint is 249.25 mm, the total width of the torso at the shoulder joint is 310 mm, and the cylinder

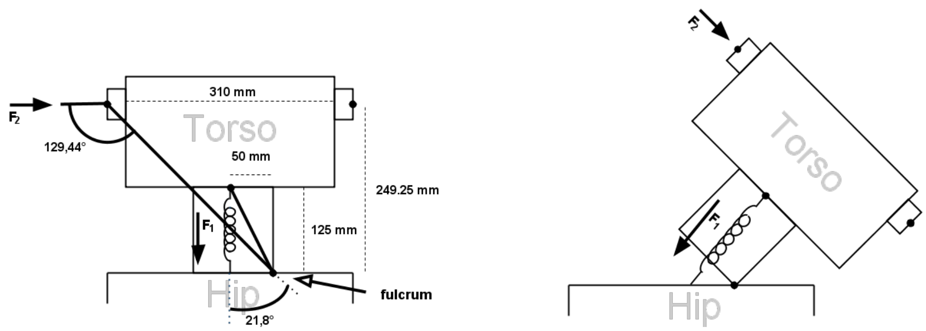

Fig. 8. Analysis of the effect of the ground reaction force. The cylinder creates a fulcrum on the hip frame and the torso works as a lever against the spring. At the shoulder a joint a force of 12.8 N (F_2) is needed to expand the spring pulling at the center of the torso with an initial force of 64 N (F_1)

has a diameter of 100 mm. When pushing the torso by the shoulder joint, the cylinder creates a fulcrum on the hip 50 mm away from the center. The spring has an angle of attack of approximately 21.8° and the ground reaction force has an angle of attack of approximately 129.44°. Taking the initial spring load of 64 N and the leverage effect into account, this results in a force of approximately 12.85 N that is needed to start extending the spring and moving the torso out of its place. This calculation neglects the effects of the arm padding and the yield of the rubber struts in the shoulder joint. Since the initial force is relatively small, we expect that the torso is moved out of place each time the goalie falls to the ground. After the impact forces are absorbed, the spring pulls the torso back and the cylinder snaps back onto the ring. If the forces are too strong or attack the torso at unexpected points and angles, it can happen that the cylinder completely detaches from the ring and the torso "dangles" off the hip. In this case, the spring is unable to pull the torso back and it has to be put back in place manually.

6 Experimental Results

Even though the soccer fields in the Humanoid Soccer league still provide structured and well defined environments, the real performance of the robots during the competitions can be regarded as a much more realistic benchmark than lab experiments. Our team NimbRo has a successful history in the Humanoid Soccer league and this can partially be accounted to the performance of our goalies. Our TeenSize goalie Bodo performed well in Suzhou in the year 2008 and successfully defended the goal several times with an earlier version of our goalie motion, as can be seen in Figure 5. We have applied our improved goalie motion as presented in this paper to our KidSize and TeenSize models and both were performing well in the German Open 2009 in Hannover and the RoboCup 2009 in Graz, where they successfully blocked several balls and prevented the opponent team from scoring. In particular in the TeenSize Dribble & Kick finals this was decisive for winning the game 2:0 against the reliable striker of the Japanese Team CIT Brains.

7 Conclusions

We have presented our solution for a robot soccer goalie motion, which is an intentional fall in a targeted direction. The intention of the fall makes it possible to execute the motion every time the same way and to prepare a pose that reliably directs the impact of the ground forces into mechanical precautions. The diving motion was developed in a simulation first and then optimized on real hardware. We reached a falling duration of approximately 0.64 seconds with our KidSize robots. The diving motion is triggered by a decision based on two factors: the point of contact and the time of contact of the ball with the lateral plane relative to the goalie. To avoid incapacitation and possible damage, the goalie only dives if it is absolutely necessary to block the ball. Our strategy to avoid damage consists mostly of mechanical solutions. The main protection of the KidSize goalie is a padded upper arm, while the goalie motion makes sure that this is the place of the first ground contact. Our TeenSize robot is protected by a flexible shoulder joint and a pull linkage in the hip held in place by a strong spring. Additionally, for both robots the joints are relaxed shortly before ground impact. None of our soccer goalies sustained any damage during games in the past years.

In future work we are planning to further automate the training process of the goalie decision and to investigate possibilities how a robot can learn to optimize the diving motion.

Acknowledgments

Funding for the project is provided by Deutsche Forschungsgemeinschaft (German Research Foundation, DFG) under grants BE 2556/2-2,/4.

References

1. J. Baltes, S. McGrath, J. Anderson, *The Use of Gyroscope Feedback in the Control of the Walking Gaits for a Small Humanoid Robot*, RobuCup 2004, Lecture Notes in Computer Science vol. 3276, pp. 628-635
2. O. Hhn, J. Gacnik, W. Gerth, *Detection and Classification of Posture Instabilities of Bipedal Robots*, Proceedings of the 8th International Conference on Climbing and Walking Robots and the Support Technologies for Mobile Machines - CLAWAR 2005, London, England, pp. 409-416
3. R. Renner, S. Behnke, *Instability Detection and Fall Avoidance for a Humanoid Using Attitude Sensors and Reflexes*, Proc. 2006 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, 2006, pp. 2967-2973

4. M. Morisawa, S. Kajita, K. Harada, K. Fujiwara, K.K. Fumio Kanehiro, H. Hirukawa, *Emergency Stop Algorithm for Walking Humanoid Robots*, Proc. 2005 IEEE/RSJ Int. Conf. in Intelligent Robots and Systems, 2005, pp. 2109 - 2115
5. J. Pratt, J. Carff, S. Drakunov, A. Goswami, *Capture Point: A Step toward Humanoid Push Recovery*, Proc. 2006 IEEE-RAS Int. Conf. on Humanoid Robots, 2006, pp. 200-207
6. S.-K. Yun, A. Goswami, Y. Sakagami, *Safe Fall: Humanoid robot fall direction change through intelligent stepping and inertia shaping*, ICRA 2009
7. K. Fujiwara, Kanehiro F., S. Kajita, K. Yokoi, H. Saito, K. Harada, K. Kaneko, H. Hirukawa. *The first human-size humanoid that can fall over safely and stand-up again*, 2003, In IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS), pp. 19201916
8. K. Fujiwara, F. Kanehiro, H. Saito, S. Kajita, K. Harada, H. Hirukawa, *Falling motion control of a humanoid robot trained by virtual supplementary tests*, 2004, In IEEE Int. Conf. on Robotics and Automation (ICRA), pp. 10771082
9. K. Fujiwara, F. Kanehiro, S. Kajita, H. Hirukawa, *Safe knee landing of a human-size humanoid robot while falling forward*, 2004, In IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS), pp. 503508
10. K. Fujiwara, F. Kanehiro, S. Kajita, K. Kaneko, K. Yokoi, and H. Hirukawa, *UKEMI: Falling motion control to minimize damage to biped humanoid robot*, 2002, In IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS), pp. 25212526
11. K. Fujiwara, S. Kajita, K. Harada, K. Kaneko, M. Morisawa, F. Kanehiro, S. Nakaoka, S. Harada, H. Hirukawa, *Towards an optimal falling motion for a humanoid robot*, 2006, In Proc. of the Int. Conf. on Humanoid Robots (Humanoids), pp. 524529
12. K. Fujiwara, S. Kajita, K. Harada, K. Kaneko, M. Morisawa, F. Kanehiro, S. Nakaoka, S. Harada, H. Hirukawa, *An optimal planning of falling motions of a humanoid robot*, 2007, In IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS), pp. 524529
13. K. Ogata, K. Terada, Y. Kuniyoshi, *Real-time selection and generation of fall damagae reduction actions for humanoid robots*, 2008, In Proc. of the Int. Conf. on Humanoid Robots (Humanoids), pp. 233-238
14. K. Ogata, K. Terada, Y. Kuniyoshi, *Falling motion control for humanoid robots while walking*, In IEEE-RAS 7th Int. Conf. on Humanoid Robots (Humanoids), Pittsburgh, 2007
15. J. Ruiz del Solar, R. Palma Amestoy, R. Marchant, I. Parra-Tsunekawa, P. Zegers, *Learning to fall: Designing low damage fall sequences for humanoid soccer robots*, 2009, Robotics and Autonomous Systems, pp. 796-807
16. C. Chang, C. Lin, *LIBSVM: a library for support vector machines*, 2001, `http://www.csie.ntu.edu.tw/~cjlin/libsvm`