Building of a Heterogeneous Segway Soccer Team Towards a Peer-To-Peer Human Robot Team

Yang Gu, Brenna Argall, Brett Browning and Manuela Veloso
School of Computer Science
Carnegie Mellon University
5000 Forbes Ave, Pittsburgh, PA 15213
{guyang, bargall, brettb, mmv}@cs.cmu.edu

Abstract

Robotic soccer is an adversarial multi-agent research domain, in which issues of perception, multi-agent coordination and team strategy are explored. One area of interest investigates heterogeneous teams of humans and robots, where the teammates must coordinate not as master and slave, but as equal participants. We research this peer-to-peer question within the domain of Segway soccer, where teams of humans riding Segway HTs and robotic Segway RMPs coordinate together in competition against other human-robot teams. Beyond the task of physically enabling these robots to play soccer, a key issue in the development of such a heterogeneous team is determining the balance between human and robot player. The first ever Segway soccer competition occurred at the 2005 RoboCup US Open, where demonstrations were held between Carnegie Mellon University (CMU) and the Neurosciences Institute (NSI). Through the execution of these soccer demonstrations, many of the challenges associated with maintaining equality within a peer-to-peer game were revealed. This paper chronicles our experience within the Segway soccer demonstrations at the 2005 US Open, and imparts our interpretation and analysis regarding what is needed to better attain this goal of teammate equality within the peer-to-peer research domain. We begin with an explanation of the motivations behind the Segway soccer and peer-to-peer research, providing details of the game rules and flow. We then describe our approach to the building of a heterogeneous Segway soccer team, in which we developed a robot-dominated soccer strategy. Robot decision making was autonomous, and the human player reacted to the robot's chosen actions. Our analysis of the experience at the US Open is presented, giving regard to both research challenges as well as difficulties in the physical execution of a Segway soccer game. We evaluate the strengths and weaknesses of our robot-driven approach within the context of game performance, as well as in contrast to the human-driven approach of our opponent team from NSI. While each team displayed either a strong bias towards the human or the robot, the intent of these peer-to-peer games is in fact teammate equality. We conclude with thoughts on the direction of future research within the Segway soccer domain.

Introduction

There has been considerable research into both human-robot interaction (Nicolescu & Mataric 2003), and multi-agent teams (Dias & Stentz 2002). Additionally, since the inception of RoboCup robot soccer (Asada *et al.* 2003), there has been considerable research into robot teams operating in adversarial environments. To our knowledge, however, there has been no work yet that combines these attributes; namely, to examine human-robot interaction within an adversarial, multi-robot setting where humans and robots are team members with similar capabilities and no clear role hierarchy (Browning, Xu, & Veloso 2004; Searock, Browning, & Veloso 2004; Browning *et al.* 2004). Segway soccer is such a domain, where human and robot are teammates, both running on the Segway platform and thus uniform in physical capabilities. In such a human-robot team, how should we define the relationship between the human and the robot within a teammate framework? We can imagine two extremes in terms of the robot autonomy. One is a fully autonomous robot, without any interaction with the human player. The other is a tele-operated robot without any autonomy at all. Of course, neither extreme condition is desirable.

Our intent in these peer-to-peer games is a truly equal relationship between teammates, where soccer plays are jointly devised and executed. A distinct effort must be made, therefore, to match human and robot capabilities, so that neither dominates the actions of the team or their teammate. From this equalization baseline then, true peer-to-peer coordination may be investigated. The creation of this baseline most likely involves limiting the human player, as they are currently more capable at on-field decision making than their robotic teammate. Specifically within the realm of teammates advising one another, human teammates run the risk of giving advice so specific that the robot is essentially tele-operated by verbal commands.

The first ever Segway soccer competition recently occurred at the 2005 RoboCup US Open, hosted by the Georgia Institute of Technology in Atlanta. The two participating teams were Carnegie Mellon University (CMU) and the Neurosciences Institute (NSI). During the game, CMU's initial strategy was so focused in thrust on robot autonomy that it placed too little importance on the human player. The result was a lack of team performance, as the robot was in reality not a strong enough player to carry that much team dependence. In contrast, NSI was able to coordinate well and accomplish tasks as a team, but at the expense of minimal robot decision making during the game.

Motivated by the observations from US Open games, we propose a peer-to-peer human robot team as our goal. Peer-to-peer means each player in the team are equally autonomous to accomplish tasks. A peer-to-peer human-robot team has no central commander, but rather each player both advising and receiving advice from their teammates, and game strategies being jointly decided upon and executed.

The format of the paper is as follows. In the next section, we describe the specifics of Segway Soccer; its goals, challenges, game flow and rules. Then we describe our development of a soccer playing Segway RMP. Following this, we describe our experience at the 2005 US Open and analyze the difference between our approach and NSI approach. We then focus on our thoughts on future human-robot games, leading to our conclusions.

Segway Soccer

Segway soccer is a game that requires mixed teams of humans and robots to cooperate to achieve the maximum reward in an adversarial task. To investigate peer-to-peer interactions, and thus ensure interesting cooperation, both humans and robots are equipped with similar capabilities. We achieve this difficult task by requiring that both humans and robots use the same drive platform, specifically the Segway platform developed by Segway LLC (Nguyen *et al.* 2004). Our goal is to create a task that requires advanced robot intelligence, combined with robust human-robot interaction skills. We hope to extend the powerful aspects of RoboCup-competition, an adversarial domain requiring fast decisions, and a well understood task - to incorporate human-robot interaction. The need for this new domain lies in the lack of study for human-robot interaction where decisions must be made quickly. As robots become more integrated into society, they will inevitably have to interact with humans and/or legacy robots in complex tasks. For some of these tasks, decisions may need to be made quickly and roles of both humans and robots may not be clearly defined a priori. In the following paragraphs, we outline the game of Segway Soccer, particularly focusing on the modifications to standard soccer rules necessary to accommodate the presence of both humans and robots on the field.

Goals and Challenges

The rules of Segway Soccer are a combination of soccer and Ultimate Frisbee [1]. The objective of the game is to score the most goals by kicked soccer ball. Adjustments to soccer rules were necessary, however, to take into account the mixture of humans and robots on the field. Most importantly, consideration was given to the size and weight of the robot, as the Segway RMP carries considerable power as a robotic platform. For safety, so that robots and humans will not contest each other for ball possession, a player in possession of the ball has a 1.0m radius in which to reposition and pass to a teammate. Furthermore, a mixed team cannot officially score unless both a robot and a human interact with the ball on the way to the goal. A passing sequence is therefore required prior to a goal shot, thus ensuring the collective involvement of both the robot and human teammate in a goal score.

[1] Rules for Ultimate Frisbee can be found at http://www.upa.org

Communications

There is no restriction on audible (speakers and microphones) communications between teammates (humans or robots). Wireless communication is allowed only between team members, and not to any off-field computers. In the spirit of RoboCup, the robots must be autonomous and make their own decisions, and thus communicate with human players only for reasons of team coordination. Some level of commands may be given to the RMP (such as waypoints, or general directions on the field as to where to play or go), but direct teleoperation of the robot is not allowed.

Game Flow

A game consists of three parts, a first half, a break, and a second half (Browning *et al.* 2005). Kickoffs occur at the start of the game or after a goal, at which time the ball is placed at the goal spot on the defensive side of the team with possession. Afterwards, players gain possession based on proximity to the ball when it is “free” or whenever the other team scores a goal. Once a player obtains possession opponents are not allowed to contest the ball, and the ball must be passed - it may not be ’dribbled’ - for the team to maintain possession. A time limit requires the ball be passed else possession be overturned.

Humans are only allowed to kick the ball with the Segway HT platform and not with any part of their bodies. To prevent double teaming, only one defender is allowed to be within 2.0m of the player currently in possession of the ball. Until the robots become more proficient, humans are not allowed to mark the opponent robot.

Upon the scoring of a goal, the game is immediately halted and then restarted from the kickoff position with a turnover in possession. Goals are only awarded when both the robot and human participate in a given play by either shooting the goal or passing to their teammate. In the original rules, no restrictions are placed on which teammate is allowed to score a goal.

Our Development of a Soccer Playing Segway RMP

Development of a robotic soccer player requires a control architecture for the execution of soccer actions. This control architecture is dependent upon the robot’s current actions, as well as its perceived state of the world. Our robot has been augmented with additional sensors for the acquisition of this world belief state, and processes its output to form an internal world model. Observation of the world and interpretation of its state is the motivating force behind the development of soccer play on our robots.

In this section we begin by describing the implementation of the control architecture for soccer play on our robots, which was accomplished via hierarchical state machines which we call skills. Within the context an example skill, we explain the structure of, and actions taken by, these state machines. Our offensive and defensive strategies, as they were first implemented for the 2005 RoboCup US Open, are described. For the realization of such actions, our robots were augmented with additional manipulators and sensors, about which details are provided. Lastly explained is the use of sensor input to update our belief of the world state, specifically with respect to vision processing and vision object tracking.

Soccer Play: Control by Skills

At the highest level, a soccer playing robot must choose to act, given the state of the world and the team goal for that state. Our robot is entirely autonomous in this process, and thus makes all on-field decisions independently of its human teammate. For the actual decision process, a control structure was implemented in which finite state machines organize into a hierarchy. We call these state machines skills. Skills process information from the believed world state and then generate action commands, either by execution of the action within the state or by calling another skill as a subskill. It is from this calling of sub-skills that a hierarchy is formed.

The state of the world (for example, the ball location being unknown) is perceived by the robot from external and internal sensors. Both interpreted perception of the world and the action being taken by the robot within that world (for example, the ball location being unknown and the robot performing a search for it) define a state within a skill. State transitions are controlled by a change in either of these defining factors (for example, the ball location is now known and the robot terminates its search), or by a timeout on the state.

In this section we will outline, within the context of an example skill, the cycling through of one of these state machines. Included in the skill outline will be descriptions of the particulars of any executed action, as well as that action's use of our added manipulators and sensors. We begin with an example skill, before going in to strategy details.

An Example Skill: Its States and Actions An example high level skill is CatchKickTo (Fig. 1). The overall goal of CatchKickTo is to receive a pass from the teammate, and then kick the ball to a specified target (teammate or goal). The skill consists of two states, each of which calls a subskill.

The sub-skill called by the first state of CatchKickTo is Catch. The Catch skill begins with a search state, looking for either a ball or teammate. Most basic to our skill building are the vision searches: every skill is dependent upon the detection of a vision object, and thus also a vision search. Vision is the primary source of world information on our robots. For the acquisition of visual data, two cameras were added to the robot (Fig. 2, the stock Segway RMP has no visual sensing capabilities).

The sub-skill Catch then aims at the ball in preparation to receive the pass, and tries to capture the ball once it is close enough, by using the sub-skill Grab. Both running up

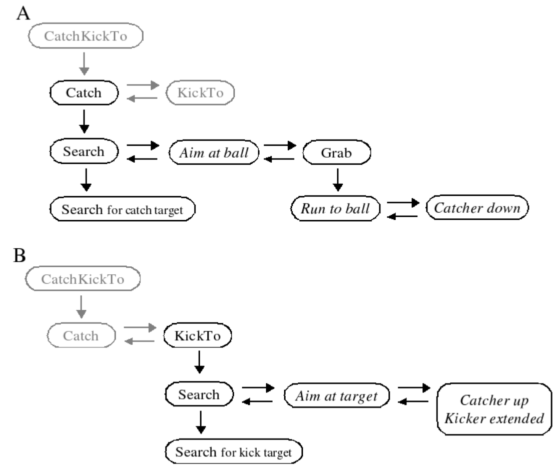

Figure 1: (A) The first state of the skill CatchKickTo calls the sub-skill Catch. Catch calls a search sub-skill, performs the action of aiming at the ball and finally calls the subskill Grab, which performs the actions of going near the ball and putting down the catcher. (B) The second state of the skill CatchKickTo calls the sub-skill KickTo. KickTo calls a search sub-skill and performs the actions of aiming at the kick target and kicking.

to the ball, and turning to aim towards the pass, require the ability to command robot motion. Actions were developed to control the motion of the Segway RMP, to which we are able to command velocity and turn angle. Our developed actions include the ability to send the robot to a given location by the specification of x, y coordinates in global or relative space, with or without obstacle avoidance (Bruce & Veloso 2002), as well as to have the robot turn in place. Once the ball is within range of the robot, its detection and capture is handled by our added infrared (IR) sensors and catcher. The two IR sensors are located near the base of the robot and are intended for ball detection, while the catcher is intended to trap the ball near the robot. Upon IR sensor firing, the catcher is commanded down and the ball is now within the robot's possession.

The second sub-skill called by CatchKickTo, is KickTo. After finding its kick target of either the teammate or the goal, KickTo calls the sub-skill AimKick. In AimKick, the robot first rotates to align itself with the identified kick target. Once aligned, the catcher is lifted and the actual kick performed. The kick is executed either via playback of a prerecorded bucking motion, or by the extension of our added pneumatic kicking plate.

Our Approach to Team Coordination: Robot Driven

The combination of skills, both simple and complex, into a large state machine constitutes our soccer game play. Soccer is a team sport, and therefore the building of our game strategy required not only execution of this large state machine, but also coordination with our teammate, the human player. Our approach was both robot and research driven; that is,

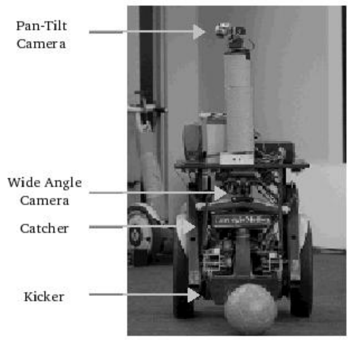

Figure 2: Our Segway RMP soccer robot equipped with a kicker, a catcher, infrared sensors, a wide-angle camera, and a camera mounted on a custom pan-tilt unit.

given our interest in the autonomy of the robot, our team strategy was a robot controlled one. There was no communication occurring from the human to the robot, and communication from robot to human was minimal. The robot would at times speak to cue the human as to its current state in the skill; for example when aiming at the ball during Catch, it would say “pass to me”. Thus all robot-to-human coordination, and the majority of human-to-robot coordination, was based upon observation alone. For example, the robot would choose to pass the ball to its teammate only after visually identifying them. As a direct result of this, all robot decision making during the demonstrations was autonomous.

Offense While in preparation for the 2005 RoboCup US Open, the initial human participation in our offensive strategy was dependent upon teammate vision almost exclusively. That is, the robot did not presume their teammate to be in any given position, nor to be taking any sort of action, besides waiting to receive the robot’s pass.

At a kickoff for our side, the robot first decided whether to run up and grab the ball, or to wait and receive a pass. This decision was dependent upon the teammate being seen, and if they were seen, whether the teammate or robot were closer to the ball. The human teammate was not searched for explicitly, but the ball was, and the teammate was often found during this ball search. The human teammate responded to the actions of the robot; that is, they waited to see if the robot was running up to or aiming at the ball, and then react accordingly. The kickoff pseudo-code runs as follows:

```
if ball seen
   if teammate seen
      if teammate closer
         Wait to receive pass.
      else
         Grab ball.
      end if
   else
      Grab ball.
   end if
else
   Search for ball.
end if
```

After kickoff, our robot’s offensive strategy presumed only to receive a pass, and therefore never ran up and grabbed the ball, as an open ball, able to be grabbed, would imply our side having lost the ball, and therefore also imply a possession turn over. Once having received the pass and being in possession of the ball, the robot then decided whether to pass to its teammate or shoot on the goal. Again, this behavior was dependent upon teammate detection, and additionally on goal detection and goal distance from the robot. If neither were seen, the robot would continue searching up to the possession timeout, at which point it would just kick the ball away from itself. This last case, however, never happened during the actual US Open demonstrations, as the field was small enough for the goals to be detected from any point on the field. After passing to the teammate, the robot would then position itself a set distance down field, towards the opponent’s goal. The game play pseudo-code is here presented:

```
if have ball
   if teammate seen
      if goal seen
         if goal distance ”¡” thresh
            Shoot on the goal.
         else
            Pass to teammate.
         end if
      else
         Pass to teammate.
      end if
   else
      Random kick after timeout.
   end if
else
   Position downfield.
end if
```

This dependence upon teammate detection as a gateway into the execution of the majority of our offensive soccer play proved to be a crutch during our actual US Open demonstrations. Our offensive strategy, therefore, was modified to presume more about our teammate’s actions and thus rely less heavily on their actual detection (for further details, please refer to section 5.2.2).

Defense Our initial defensive strategy relied more heavily on the human player. The robot positioned itself inside of our goal with the intent of catching attempted shots on the goal. Any defensive actions besides goal keeping - such as attempting to intercept or block opponent passes - were performed by the human player.

Constraints on field size at the US Open seriously restricted the motion of our robots, and thus made consistent positioning inside the goal infeasible. Our defensive strategy was therefore likewise modified, to take the robot out of the goal and mark the ball in attempts to gain possession of it (for further details, please refer to section 5.2.3).

Perception

Perception is where the robot autonomy comes from. In this section, we briefly introduce the two key components, vision and tracking in our large system. Based on vision, we construct our world model to further implement robot autonomy. Since the objects we have interest in are not always visible, we need tracking to estimate their position and velocity consistently.

Vision In our work with the Segway RMP platform (Nguyen *et al.* 2004), we have been exploring techniques to enable a vision-centric robot to be able to play soccer in outdoor environments where illuminate is variable (Browning *et al.* 2004; Browning & Veloso 2005). Furthermore, as the task is adversarial and highly dynamic, the combination of robot speed and ball speed means that it is essential that the vision algorithms be both robust, and extremely efficient. Indeed, only a fraction of the CPU resources can be devoted to vision processing as the remainder of the CPU resources must be used for cognition in order to get low-latency robot behaviors.

We have developed a new technique for fast color-object recognition that is suitable for use in robot platforms like the Segway. Its key feature is that it is able to adapt its segmentation process to different lighting conditions. The segmentation technique is motivated by the observation that for most of the domains of interest here changes in illumination lead to small changes in color value and that these changes are relatively uniform across all colors. In other words, with modern cameras with automatic shutters and gain control red pixels may vary in color but will stay in the same region of color space. Therefore, we propose that if pixel classification thresholds are able to adapt by small amounts, it should become possible to robustly classify pixel colors across moderate changes in illumination. Our goal is to achieve such a robust, adaptive system but without significantly increasing computational requirements. For more details of this approach, see our paper (Browning & Veloso 2005).

The key idea is to use a soft labeling of pixel class, followed by a hard decision using an adaptive threshold. The combination of soft-labeling and adaptive thresholding provides the plasticity for lighting variation. Following this, connected pixels can be conglomerated using a connected component analysis. Objects are detected and recognized by searching for nearby regions that match priori models, with soft-comparisons to account for variations in shape, size, and missing features. This new technique requires only moderate additional computational resources beyond existing fast color vision algorithms.

Tracking Tracking in essence consists of using sensory information combined with a motion model to estimate the position of a moving object. Tracking efficiency completely depends on the accuracy of the motion model and of the sensory information (Y.Bar-Shalom 2001). When tracking is performed by a robot executing specific tasks acting over the object being tracked, such as a Segway RMP soccer robot grabbing and kicking a ball, the motion model of the object becomes complex, and dependent on the robot's actions (C. Kwok 2004). A single motion model is not exact enough to describe the complexity of the motion due to the interactions between the robot and the ball. We therefore use a tactic-based multiple model approach to model the ball motion. Explicitly, we use the following three single models.

- *Free-Ball*. The ball is not moving at all or moving straight with a constant speed decay d which depends on the environment surface.
- *Grabbed-Ball*. The ball is grabbed by the robot's catcher.
- *Kicked-Ball*. The ball is kicked therefore its velocity is equal to a predefined initial speed plus the noise.

Next, a model index m determines the present single model being used ($m = 1, 2, 3$ for the above three single models respectively). We need to decide how to transit between each models, which is done by a tactic based approach. We assume that the model index, m_k, conditioned on the previous tactic executed t_{k-1}, and other useful information v_k (such as ball state $\mathbf{x}_{k-1}$, infrared measurement s_k or the combination of two or more variables), is governed by an underlying Markov process, such that, the conditioning parameter can branch at the next time-step with probability

$$p(m_k = i | m_{k-1} = j, t_{k-1}, v_k) = h_{i,j} \qquad (1)$$

where $i, j = 1, \cdots, N_m$.

Finally, we use particle filtering to track the motion model m and the ball state b (S.Arulampalam *et al.* 2002). Particle filter maintains the belief state at time k as a set of particles $p_k^{(1)}, p_k^{(2)}, \cdots, p_k^{(N_s)}$, where each $p_k^{(i)}$ is a full instantiation of the tracked variables $\mathbf{P}_k = \{p_k^{(i)}, w_k^{(i)}\}$, $w_k^{(i)}$ is the weight of particle $p_k^{(i)}$ and N_s is the number of particles. In our case, $p_k^{(i)} = \langle b_k^{(i)}, m_k^{(i)} \rangle$.

We use the Sample Importance Resampling (SIR) algorithm to update the state estimates (A. Doucet 2001). The sampling algorithm is as follows:

$$[\{b_k^{(i)}, m_k^{(i)}, w_k^{(i)}\}_{i=1}^{N_s}] = \textbf{\textit{SIR}}[\{b_{k-1}^{(i)}, m_{k-1}^{(i)}, w_{k-1}^{(i)}\}_{i=1}^{N_s}, z_k, s_k, t_{k-1}]$$

01 **for** $i = 1 : N_s$
02 draw $m_k^{(i)} \sim p(m_k | m_{k-1}^{(i)}, b_{k-1}^{(i)}, s_k, t_{k-1})$.
03 draw $b_k^{(i)} \sim p(b_k | m_k^{(i)}, b_{k-1}^{(i)})$.
04 set $w_k^{(i)} = p(z_k | b_k^{(i)})$
05 **end for**
06 Calculate total weight: $w = \sum [\{w_k^i\}_{i=1}^{N_s}]$
07 **for** $i = 1 : N_s$
08 Normalize: $w_k^i = w_k^i / w$
09 **end for**
10 ***Resample***.

The inputs of the algorithm are samples drawn from the previous posterior $\langle b_{k-1}^{(i)}, m_{k-1}^{(i)}, w_{k-1}^{(i)} \rangle$, the present vision and infrared sensory measurement z_k, s_k, and the tactic t_{k-1}. The outputs are the updated weighted samples $\langle b_k^{(i)}, m_k^{(i)}, w_k^{(i)} \rangle$. In the sampling algorithm, first, a new ball motion model index, $m_k^{(i)}$, is sampled at line 02. Then given the model index, and previous ball state, a new ball state is sampled at line 03. The importance weight of each sample is given by the likelihood of the vision measurement given the predicted new ball state at line 04. Finally, each weight is normalized and the samples are resampled. Then we can estimate the ball state based on the mean of all the $b_k^{(i)}$. For more details of the tracking, see our paper (Gu 2005).

Opponent Approach to Team Coordination: Human Driven

Our opponent Segway team was developed by the Neurosciences Institute in San Diego, CA (http://vesicle.nsi.edu/nomad/segway/). The sensing capabilities on the NSI robot included a pan-tilt CCD camera, a forward facing SICK laser, and both forward and backward facing IR sensors. Additional manipulators were a solenoid powered catcher and kicker to grab and kick the ball, as well as a ball capture mechanism. This mechanism to our knowledge consisted of a skirt of tubing intended to trap a possibly out of sight ball, which the robot would then sense and turn towards until the ball rested within the catcher. Their human Segway was likewise outfitted with a similar catcher and kicker, and additionally a headset through which the human player could communicate with its robot teammate. Robot movement was guided by neural simulation, where sensor input generated motor signals to command velocity and turn angle at the Segway wheels. In contrast to our robot-dominated approach to peer-to-peer team coordination, NSI developed a human-dominated game strategy. Their human player performed the majority of the decision making on the field, and then informed their robot, by voice, of the chosen action. To our knowledge, the actions spoken by the human player to the robot included whether and when to shoot on the goal or pass to the teammate, as well as guidance on location choice for field positioning.

The 2005 US Open Experience

Five Segway soccer demonstrations were played between Carnegie Mellon University and the Neurosciences Institute at the 2005 US Open.

Logistics Difficulties

The actual execution of multiple Segway soccer demonstrations made evident several issues with the game implementation, both as a result of the stated game rules as well as the setup of the physical space. In this section we describe our observations regarding what these issues were, as well as our interpretation of their cause.

Robot Movement In an ideal peer-to-peer game, equal amounts of teammate mobility would be shown. Such equality is necessary not only in the interest of normalizing capabilities, but also because a bias in mobility will undoubtedly lead to a bias in field performance and participation. The US Open demonstrations, however, overall displayed a marked lack of robot positioning. We believe the cause of this reduced mobility to be twofold.

The first culprit which constrained robot movement was field dimension. Due to size constraints at the US Open venue, the field occupied approximately a quarter of the area as was originally stated in the rules, being halved in each dimension. The second culprit confounding robot movement

Figure 3: The field at US Open 05 was too tight to pass and position easily.

was the safety distance (of 1.0m) required between players, which by and large was respected by the robots. That the human players were able to maneuver more easily was due largely to their disregard for, and the difficulty of referee enforcement of, this distance rule. In the early demonstrations, CMU navigation was particularly conservative, and therefore the robot practically immobile. Additionally, this rule was interpreted differently by each team; the 1.0m as a radius was defined by CMU as extending from a point particle centered on the robot, and by NSI as extending from the outer perimeter of the robot.

The reduction in field size, compounded with the distance restriction between players, so congested the field that robots frequently were unable to position themselves (Fig. 3). This lack of positioning had the immediate effect of a reduction also in passing between teammates, where often the more mobile human player would execute only the minimum requirement of one pass to its robot teammate before shooting on the goal. The small size of the field additionally encouraged such behavior by placing the goal within reasonable shooting distance from most points on the field.

Robot Participation in Passing Equality in how each teammate participates in a pass is likewise required of a true peer-to-peer game. The human and Segway teammates are not completely normalized across their wheeled platforms, but rather are still divided by the very basic reality that the bodies above the platforms are physically different. Such a distinction logically might result in a difference in action execution on the playing field. At the US Open, we observed this within the context of teammate passing.

While ball deflection is common within human soccer games, it is also unlikely that a human player would attempt to deflect the ball off of an unsuspecting teammate. Such consideration is no longer necessary, however, when that teammate is no longer human. These deflections, as observed in use by NSI, were at times again caught by the human player. It is possible that the robot knew of these passes and intentionally missed them to allow for more time, or attempted to catch the ball but failed, but it is also possible that the robot was not aware such passes occurred at all. In the spirit of peer-to-peer games, the robot should be an active and aware participant in any coordinated actions, and, properly constructed, the game rules should enforce this spirit.

However, to determine robot awareness explicitly, and not just intuitively, is a difficult and situation dependent task.

Another question presented by the experience was what classifies an acceptable pass. Within human soccer games ball possession is not guaranteed for a set radius around a player, and so the ball may be more aggressively contested than in Segway soccer. Within Ultimate Frisbee, a pass is considered successful only if the receiving player actually catches the frisbee, and is enforced by requiring it never touch the ground. In Segway soccer, it is possible a pass might be considered valid though it actually remains untouched by the player, since the receiving player gains possession should the ball come within the 1.0m safety distance. Judging the validity of such a pass is unable to be helped by the rules of Ultimate Frisbee, since ofen the ball has never left the ground in the first place. As such, what fairly defines a received pass requires further investigation.

Robot Goal Scoring Peer-to-peer teammates would be expected to attempt shots on the goal equally across platform; that is, delegation to a specific soccer role might influence a player's shooting frequency, but whether they are human or robot should not. By the completion of the third demonstration, however, a goal had yet to be scored by a robot. Beyond positioning difficulties allotting robots fewer opportunities to score in the first place, any attempted robot shots on the goal were blocked by the human opponent. Considering the game to be too human-dominated, the teams agreed to a rule addition to restrict the human players: a human was no longer allowed to score a goal. While the instantiation of this restriction resulted in many fewer attempted and successful goals, that all were scored by robots increased their participation in the game dramatically.

The CMU Experience

In the following section we describe the experience of our team particularly at the US Open, both with respect to the afore mentioned logistical difficulties and our opponent team, as well as our resultant adaptation in offensive and defensive strategies.

Initial Analysis The most obvious failing the robot displayed when executing our initial strategies (as described in section 3.2) was appearing to be in a constant state of search.

While the goals were large and nearly always detected, the ball and teammate were often occluded and therefore not. By making no offensive assumptions about its teammate's behavior, the robot was dependent upon teammate detection as a gateway into transitioning to execute the remainder of a play. The folly in this was that often the remainder of a play might have still been successful even without teammate detection. For example, if instead of continuing to search for the teammate the robot had just quickly kicked the ball forward, the human teammate, able to easily detect the robot and therefore likely already positioned appropriately, would have often been able to receive the pass anyhow. The robot's defensive goal blocking required ball detection, but the reality of a full ball search often had the robot looking away from a partially occluded ball when a shot was attempted.

Figure 4: The left figure shows our Segway robot holding the ball and turning to search for its teammate. The right figure shows our Segway robot positioning to receive a pass, and being marked by the NSI robot.

Even when the ball was detected, the robot's interception capabilities were generally slower than an attempted shot on the goal.

Additional problems with the use of the robot as goalkeeper was the observed difficulty in robot positioning, due to the reduced field size. In the frequent case of the robot being unable to position properly, an eventual timeout would cause the robot to begin defending the goal even if it was not within the goal. That is, the robot would search for the ball and in the event of ball detection would attempt interception, but only allowing itself to cover minimal distance in this interception presuming itself to be in the goal and therefore the ball to be its concern only if very near.

Evolved Offense As our offensive strategy developed, coordination with the teammate, and therefore presumptions about their actions, became stronger. At a kickoff for our side, the robot assumed their teammate to begin with the ball, and therefore was positioned advantageously to attempt a shot on the opponents' goal. However, should the robot always shoot on the goal at a kickoff, this behavior would be easily predicted, and therefore blocked, by the human opponent player. An element of randomness, therefore, was added. With a predefined but configurable probability, the robot chose to either kick on the goal or at a set angle from the goal. The human player would position themselves to receive this off-goal kick, and the robot presumed the human to be in that position. That is, the off-goal kick was not dependent upon the robot visually detecting its human teammate. This randomness in goal on- or off-shooting was used throughout the offensive play. An additional element of randomness was introduced to the actual goal shot by having the robot aim towards the left, center or right sides within the goal, with equal probability. Shooting on the goal required proper perception of the goal, and so, if necessary, the goal was searched for throughout the offensive play. Should the goal not be detected, or its distance to the robot be considered out of range, the ball was kicked to the teammate. Should neither be detected, the ball was kicked forward after a timeout.

Evolved Defense Changes to our defensive strategy brought the robot out of the goal. The robot's defensive target was now to intercept the ball. While respecting the distance minimums required between players, the robot at all times attempted to grab the ball; that is, the ball was marked, and should the ball no longer be within the distance minimum of another player, it would be grabbed. This strategy proved far more effective than the robot as goal keep. Not only did the robot often effectively position itself between the opponent players, thus obstructing an intended pass, but on occasion an opponent pass was not just blocked but actually intercepted.

Future Human-Robot Games

Each team was unaware, until the first game, of the development angle chosen by the other team; that our strategies were opposite in player dominance was not intentional, but their contrast did exemplify many of the difficulties with the development of human-robot balance within the game. CMU's initial strategy, so focused in thrust on robot autonomy, placed too little importance on the human player. The result was a lack of team performance, as our robot was in reality not a strong enough player to carry that much team dependence. In contrast, NSI was able to coordinate well and accomplish tasks as a team, but at the expense of minimal robot decision making during the game. As the intent of this research domain is true human-robot coordination, where the players are equally autonomous yet also able to accomplish tasks, it seems a balance somewhere between the two approaches must be found. Such a balance will by necessity restrict the human players initially, but as the robots become more capable, so also will interspecies equality between teammates become less artificial.

Conclusions

Soccer is an adversarial multi-agent coordination task, currently a research domain within multiple robot platforms. The intent of Segway soccer specifically is to research the concept of peer-to-peer games; that is, games in which humans and robots coordinate in soccer play as equal teammates. We have proposed these peer-to-peer human-robot teams as the goal of future human-robot games. Our approach to the development of a Segway Soccer team was robot-dominated, while the approach of our opponent NSI team was human-dominated. We have evaluated the strengths and weaknesses of each approach, as observed at the 2005 RoboCup US Open. That either approach was dominated by a single species, however, counters the intents of equality in peer-to-peer games. Our aim in further developing this domain is to encourage both robot autonomy and human involvement, thereby encouraging also the peer-to-peer ideal of balanced and interesting interactions between them within a soccer game.

Acknowledgment

This work was supported by United States Department of the Interior under Grant No. NBCH-1040007. The content of the information in this publication does not necessarily reflect the position or policy of the Defense Advanced Research Projects Agency (DARPA), US Department of Interior, US Government, and no official endorsement should be inferred.

References

A. Doucet, N. De Freitas, N. G., ed. 2001. *Sequential Monte Carlo Methods in Practice*. New York: Springer-Verlag.

Asada, M.; Sklar, E.; Kaminka, G. A.; Obst, O.; Polani, D.; Browning, B.; Bonarini, A.; Fujita, M.; Christaller, T.; Takahashi, T.; and Tadokoro, S. 2003. An overview of robocup-2002 fukuoka/busan. *AI Magazine* 24(2):21–40.

Browning, B., and Veloso, M. 2005. Real-time, adaptive color-based robot vision. In *Proceedings of IROS'05*.

Browning, B.; Rybski, P.; Searock, J.; and Veloso, M. 2004. Development of a soccer-playing dynamically-balancing mobile robot. In *Proceedings of International Conference on Robotics and Automation*.

Browning, B.; Rybski, P.; Gu, Y.; and Argall, B. 2005. Segwayrmp robot football league rules. In *Coral Publications*.

Browning, B.; Xu, L.; and Veloso, M. 2004. Skill acquisition and use for a dynamically-balancing soccer robot.

Bruce, J., and Veloso, M. 2002. Real-time randomized path planning for robot navigation. In *Proceedings of IROS-2002*. An earlier version of this paper appears in the Proceedings of the RoboCup-2002 Symposium.

C. Kwok, D. F. 2004. Map-based multiple model tracking of a moving object. *Proceedings of eight RoboCup International Symposium*.

Dias, M. B., and Stentz, A. 2002. Opportunistic optimization for market-based multirobot control. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems IROS 2002*.

Gu, Y. 2005. Tactic-based motion modelling and multi-sensor tracking. *Proceedings of Twentieth National Conference on Artificial Intelligence*.

Nguyen, H. G.; Morrell, J.; Mullens, K.; Burmeister, A.; Miles, S.; Thomas, K.; and Gage, D. W. 2004. Segway robotic mobility platform. In *SPIE Mobile Robots XVII*.

Nicolescu, M., and Mataric, M. J. 2003. Natural methods for robot task learning: Instructive demonstration, generalization and practice. In *Proceedings of the Second International Joint Conference on Autonomous Agents and Multi-Agent Systems*, 241–248.

S.Arulampalam; S.Maskell; N.Gordon; and T.Clapp. 2002. A tutorial on particle filters for on-line non-linear/non-gaussian bayesian tracking. *IEEE Transactions on Signal Processing* 50(2):174–188.

Searock, J.; Browning, B.; and Veloso, M. 2004. Turning segways into robust human-scale dynamically balanced soccer robots.

Y.Bar-Shalom, X.-R.Li, T. 2001. *Estimation with Applications to Tracking and Navigation*. John Wiley & Sons, Inc.