

The First Segway Soccer Experience:
Towards Peer-to-Peer Human-Robot Teams

Brenna Argall Yang Gu Brett Browning
 Manuela Veloso
 July 2005
 CMU-CS-05-161

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

In this paper, we focus on human-robot interaction in a team task where we identify the need for peer-to-peer (P2P) teamwork among robots and humans. We define a P2P team as one where there is no fixed hierarchy for decision making. Instead, all team members are equal participants and decision making is truly distributed. We have fully developed a P2P team within Segway Soccer, a research domain that we have introduced to explore the challenge of P2P coordination in human-robot teams in dynamic, adversarial tasks. Segway soccer builds upon RoboCup robot soccer, a research domain with standardized tasks for comparing multi-robot team strategies. We recently participated in the first Segway Soccer games between two competing teams at the 2005 RoboCup US Open. We believe these games are the first ever between two human-robot P2P teams. Based on the competition, we realized two different approaches to P2P teams. We present our robot-centric approach to P2P team coordination and contrast it to the human-centric approach of the opponent team. Based on our analysis, we recommend a set of further refinements to the Segway Soccer domain to foster more effective P2P teams.

Keywords: Segway soccer, human-robot teams

1 Introduction

Within the field of robotics, there has been considerable research into multi-robot coordination for a variety of domains and tasks (e.g., [3, 11, 13, 21, 25, 34]). In particular, there has been an ongoing effort to explore multi-robot teamwork in dynamic, adversarial tasks within a standardized domain: RoboCup robot soccer [2, 18, 27, 33] (see also <http://www.robocup.org>), which offered a pioneering multi-robot challenge. Throughout, a variety of techniques have emerged to effectively coordinate teams robots. Inevitably we need distributed approaches in which each team member is able to make decisions and communicate [12, 26, 28, 32]. An ongoing challenge in multi-robot research is how to effectively coordinate heterogeneous teams [4, 17, 19, 30]. In this paper, we focus on such teams composed of humans and robots, as opposed to only robots, specifically in dynamic tasks that require robust real-time action.

Much of the recent work on human-robot interaction [14, 16, 23, 24] focuses on the challenging problem of how to effectively communicate commands (both low and high-level) and information between a human and a robot. Indeed, there is often a well defined decision making hierarchy whereby the human is the leader. We call such an approach a human-centric team strategy. We conjecture, that in many team scenarios the real world factors of dynamics, partial observability, and communication limits, mean that such coordination may produce sub-optimal team performance. That is, the inflexibility of such a master-slave relationship may degrade the team's ability to respond to dynamic opportunities or problems as would a centralized approach to multi-robot coordination.

Our insights come from our work on developing a scenario where teams of humans and robots would face adversarial environments as peers. We have developed Segway Soccer to explore balanced, flexible, and distributed approaches to human-robot coordination, as we introduce *Peer-to-Peer* (P2P) teams. P2P teamwork is encouraged by requiring both humans and robots to have identical motion capabilities by deriving from the Segway mobility platform. Specifically, humans ride Segway Human Transports (HT) while robots are built from Segway Robot Mobility Platforms (RMP).

In this paper, we concretely define a P2P team as one where team members are equal, in the sense that each team member has equal power to make decisions and the responsibility to do so. Thus, there is no fixed decision making hierarchy, no fixed leader, and coordination is distributed. Key decisions will flow back and forth across the team as the situation dictates.

Segway soccer aims at providing a standardized research domain to ex-

plore and evaluate P2P teams. Building upon our extensive experience with different RoboCup robot soccer leagues and teams, we designed the new Segway soccer league consisting of two human-robot teams that play against each other with a precise set of soccer and ultimate-frisbee based rules [6]. As with RoboCup, the robots must be fully autonomous [5, 7, 31].

Recently, the first ever Segway soccer games were held as part of the 2005 RoboCup US Open, hosted by the Georgia Institute of Technology. We believe these were the first human-robot P2P competition with two participants, namely CMBalance'05, our team from Carnegie Mellon University (CMU) and BrainWorks, from the Neurosciences Institute (NSI) [20]. The games truly enriched our understanding of P2P teams, which we report in this paper. We realized two different approaches to P2P teams. We present our robot-centric approach to P2P team coordination and contrast it to the human-centric approach of the opponent team. Based on our analysis, we recommend a set of further refinements to the Segway Soccer domain to foster more effective P2P teams.

Concretely, the format of the paper is as follows. Section 2 describes the specifics of Segway soccer, namely its goals, challenges, game flow and rules. Section 3 briefly summarizes our development of a soccer playing Segway RMP. Section 4 focuses on our approach to the development of a Segway soccer team, including an overview of our offensive and defensive strategies, the analysis of their respective game performances, and the modifications which were made for their improvement. Following this, we describe our experience at the 2005 US Open and analyze the difference between our approach and NSI approach. We then, in Section 6, focus on our thoughts on future human-robot games and propose a revised set of Segway soccer game rules, leading to our conclusions.

2 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 [22]. 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 do-

main 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.

2.1 Goals and Challenges

The rules of Segway Soccer are a combination of soccer and Ultimate Frisbee¹. 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.

2.2 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.

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

2.3 Game Flow

A game consists of three parts, a first half, a break, and a second half [6]. 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.

2.4 The Referee Box

The referee box is a program tasked to link the referee whistle and robot understandable messages (Fig. 1). While the intent is to standardize and have all teams run the same program, as in other RoboCup leagues, for the 2005 US Open the referee boxes were developed by each team individually. Our implementation specifically requires human interaction with a GUI, where the button corresponding to the referee whistle is selected. In this way, the robot is updated with the current game state.

In the current setup, the game state is transmitted to the robotic players as follows. Before a game begins, the off-field human controlling the referee box clicks one of the “Connect” buttons (depends on the team color) to connect with the robot players on the field. Whenever a whistle is heard, we first “Stop” the robot, then click one of the game control buttons, “Kick Off”, “Throw In”, “Turn Over” or “Free Kick”, followed by “Start” to inform the robot that it is time to run again.




Figure 1: The Referee box communicates wirelessly with the robotic soccer players. It is controlled by an off-field human, via the Referee box GUI.

3 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.

3.1 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




Figure 2: (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 sub-skill 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.

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 sub-skill. 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.

3.1.1 An Example Skill: Its States and Actions

An example high level skill is CatchKickTo (Fig. 2). 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 sub-skill.

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 depen-

dent 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. 3, 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 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 [9], 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.

3.2 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, 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 identify-




Figure 3: 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.

ing them. As a direct result of this, all robot decision making during the demonstrations was autonomous.

3.2.1 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. Pseudo-code for the kickoff control structure, as well as for the game play described below, is presented in Fig. 4.

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

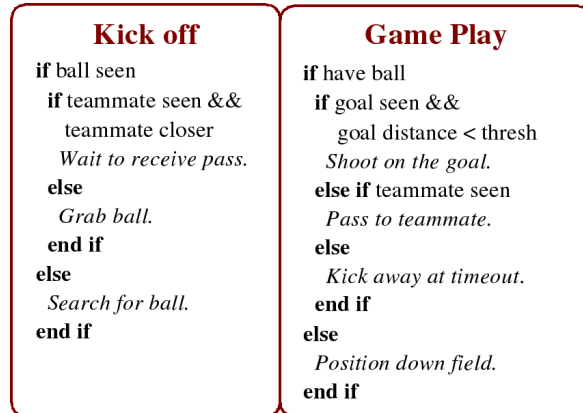


Figure 4: The psuedo-code for our kickoff (left) and game play (right) offensive soccer strategies.

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.

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).

3.2.2 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).

3.3 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.

3.3.1 Vision

In our work with the Segway RMP platform [22], we have been exploring techniques to enable a vision-centric robot to be able to play soccer in outdoor environments where illumination is variable [5, 8]. 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 [8].

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. Follow-

ing 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.

3.3.2 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 [35]. 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 [10]. 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

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

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

Finally, we use particle filtering to track the motion model m and the ball state b [29]. A particle filter maintains the belief state at time k as a set of particles $p_k^{(1)}, p_k^{(2)}, \dots, p_k^{(N_s)}$, 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 [1]. The sampling algorithm is as follows:

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

```

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 $b_k^{(i)}$. See (Fig. 5) for an example of particle filter representation of the ball position. For more details of the tracking, see our paper [15].

4 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 capa-




Figure 5: In both figures, the red and pink cross represent the two goals, and each blue dot represents the one of the particle estimation of the ball position. In the left figure, the ball is visible to the robot. In the right figure, the ball is out of sight and the particles are scattered

bilities 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.

5 The 2005 US Open Experience

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

5.1 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.

5.1.1 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 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. 6). 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.



Figure 6: The field at US Open 05 was too tight to pass and position easily.

5.1.2 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 often the ball has never left the ground in the first place. As such, what fairly defines a received pass requires further investigation.

5.1.3 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.

5.2 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.

5.2.1 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



Figure 7: 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.

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. 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.

5.2.2 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.

5.2.3 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.

6 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.

With the aim of more balanced human-robot team coordination, we propose the following rule changes [6].

6.1 Proposed Change of Rules

6.1.1 Field Size

The reduced field size at US Open 05 restricted robot movement, and thus restricted both positioning and passing.

Alternative solutions:

- A bigger indoor space.
- To take the game to an outside field.

We propose a bigger indoor field for 2006, and to go outside later.

6.1.2 Communication

In the spirit of RoboCup, coordination between teammates on the field should be equal, with neither teammate dominating the other. Similarly then should communication between teammates be equal, with neither human nor robot consistently verbally commanding the other.

Alternative solutions:

- Limit what the human is able to tell the robot, either by predefining a list of commands, or by placing a limit on the size and rate of messages.
- Require equal amounts of communication between robot and human; that is, the number of human-to-robot commands must be matched by the number of robot-to-human commands.
- Limit the human to reporting only his own information, for example “I am open” instead of “Pass to this location (me), at this time (now)”.

While requiring equal amounts of communication seems the right solution, we are uncertain as to how such a rule might be reasonably enforced.

6.1.3 Human Shooting

The issue of human domination in goal scoring was observed, and to a certain extent addressed, during the US Open. Human scoring was at that time banned, but other possible solutions exist with varying degrees of human restriction.

Alternative solutions:

- Only robots are allowed to shoot on the goal.
- Only robots are allowed to defend the goal.
- For a human goal to count, a robot must have touched the ball within the last 2 seconds (thereby retaining option for a human deflection into the goal).

Since simple enforcement of a 2 second rule seems difficult, for 2006 we propose that only robots are allowed to shoot on or defend the goal.

6.1.4 Human Marking

Restrictions on the amount of defense directed towards a robot are required if any sort of reasonable offense is expected from the current robot players. Already in existence was the rule of no double teaming (where no two players may block the same robot) and no human being allowed to mark a robot. However, a more interesting display of robot capabilities might be possible should further restrictions be implemented.

Alternative solutions:

- No interspecies marking. That is, humans mark humans, robots mark robots.
- Require the maintenance of a predefined distance between the marker and marked players.
- Enforce a timeout on marking a robot.

We support the instantiation of no interspecies marking.

6.1.5 Avoidance Distance

Humans were by far the worst offenders of the minimum distance rule, while the robots tended to both calculate and respect the avoidance distance. By the same challenge-in-calculation reasoning, as the referee is human, the

maintenance of such a distance is difficult to enforce. However, the definition of this distance, though unclear in the current rules, can be concretely stated.

Solution, with two aspects:

- The distance is defined as a radius extending from a point particle centered on the robot, since the robots as augmented for soccer vary in size from team to team.
- For safety, this distance should be increased from 1.0m, but also needs to be flexibly enforced.

6.1.6 Timeouts

The number and length of timeouts were not predefined, and were abused throughout the games.

Solution:

- A maximum of 3 timeouts, with a maximum total time 10 minutes.

6.1.7 Passing

No rule currently exists which restricts the use of a robot teammate as a springboard. Consequently, robots can “catch” the ball as humans cannot; that is, with their backs turned and the ball merely deflecting off of them. While we would like to retain the ability for deflection goals, we also want the robot to be aware of and actively involved in the action.

Solution:

- Contact with the ball must happen at the kicker for a catch to be considered valid. For example, a pass bounced off the back of the robot is no longer considered good, but a pass bounced off the kicker of either player is considered good.

6.1.8 Possession Clock

The current possession clock is too long and unclear.

Solution, with two aspects::

- The possession clock is reduced from 30 to 10 seconds.
- The clock begins to countdown when the ball touches the kicker.

6.1.9 Goal Space

We want to improve player interaction with the goals. A particular difficulty at the 2005 US Open was that the 1.0m safety distance prevented any (robot or human) player from defending their goal should an opponent be attempting a shot within 1.0m of it.

Alternative solutions:

- Instantiate a goal box, inside of which either no human player is allowed or human players have a timeout.
- Instantiate a try line instead of goal; that is, a teammate beyond the field line must catch the goal shot. (This requires at the least a larger field space, and at best an outside field.)

We find either solution acceptable for 2006.

6.1.10 Self Referee

As a suggestion, to borrow another rule from Ultimate Frisbee, we might do away with the field referee and have all players call their own fouls.

7 Conclusions

In summary, we have introduced the concept of peer-to-peer (P2P) teamwork for human-robot teams, where a P2P team is one where all teammates are equal and decisions are made in a distributed manner. We believe P2P human-robot coordination to be a powerful concept that is of great import to effective human-robot team coordination in dynamic environments. We have presented the development of Segway soccer as a standardized research domain for exploring and validating P2P coordination techniques. After participating in the first Segway soccer games at the RoboCup US Open, which we believe to be the first of its kind, we identified a number of key changes to the structure of the Segway soccer domain to further promote balanced human-robot interaction. Concretely, we have proposed rule modifications that encourage greater robot autonomy and participation in the key offensive portions of the game, and more complex human-robot interaction during general game play. We will further explore these proposed changes in our future work.

8 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

- [1] N. Gordon A. Doucet, N. De Freitas, editor. *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, New York, 2001.
- [2] Minoru Asada, Hiroaki Kitano, Itsuki Noda, and Manuela Veloso. RoboCup: Today and tomorrow – What we have have learned. *Artificial Intelligence*, 110:193–214, 1999.
- [3] T. Balch and R.C. Arkin. Behavior-based formation control for multiagent robot teams. *IEEE Transactions on Robotics and Automation*, December 1998.
- [4] Jonathan Brookshire, Sanjiv Singh, and Reid Simmons. Preliminary results in sliding autonomy for coordinated teams. In *Proceedings of The 2004 Spring Symposium Series*, March 2004.
- [5] B. Browning, P. Rybski, J. Searock, and M. Veloso. Development of a soccer-playing dynamically-balancing mobile robot. In *Proceedings of International Conference on Robotics and Automation*, May 2004.
- [6] Brett Browning, Paul Rybski, Yang Gu, and Brenna Argall. Segwayrmp robot football league rules. In *Coral Publications*, 2005.
- [7] Brett Browning, Jeremy Searock, Paul E. Rybski, and Manuela Veloso. Turning segways into soccer robots. *Industrial Robot*, 32(2):149–156, 2005.
- [8] Brett Browning and Manuela Veloso. Real-time, adaptive color-based robot vision. In *Proceedings of IROS'05*, 2005.
- [9] James Bruce and Manuela Veloso. Real-time randomized path planning for robot navigation. In *Proceedings of IROS-2002*, Switzerland, October 2002. An earlier version of this paper appears in the Proceedings of the RoboCup-2002 Symposium.

- [10] D. Fox C. Kwok. Map-based multiple model tracking of a moving object. *Proceedings of eight RoboCup International Symposium*, July 2004.
- [11] M. Bernadine Dias and Anthony Stentz. Opportunistic optimization for market-based multirobot control. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems IROS 2002*, 2002.
- [12] M Bernardine Dias, Robert Michael Zlot, Nidhi Kalra, and Anthony (Tony) Stentz. Market-based multirobot coordination: A survey and analysis. Technical Report CMU-RI-TR-05-13, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, April 2005.
- [13] Brian P. Gerkey and Maja J. Mataric. Multi-robot task allocation: analyzing the complexity and optimality of key architectures. In *Proceedings of ICRA '03, the 2003 IEEE International Conference on Robotics and Automation*, Taiwan, May 2003.
- [14] R. Gockley, A. Bruce J. Forlizzi, M. Michalowski, A. Mundell, S. Rosenthal, B. Sellner, R. Simmons, K. Snipes, A.C. Schultz, and J. Wang. Designing robots for long-term social interaction. In *In Proceedings of the Conference on Intelligent Robots and Systems (IROS)*, Edmonton, Alberta, 2005.
- [15] Yang Gu. Tactic-based motion modelling and multi-sensor tracking. *Proceedings of Twentieth National Conference on Artificial Intelligence*, 2005.
- [16] Jenks, Diftler, and Williams. Robonaut: A telepresence base astronaut's assistant. In *SPIE-2001*, Boston, November 2001.
- [17] G.A. Kaminka and I. Frenkel. Flexible teamwork in behavior-based robots. In *In Proceedings of the National Conference on Artificial Intelligence (AAAI-2005)*, 2005.
- [18] H. Kitano, M. Tambe, P. Stone, M. Veloso, S. Coradeschi, E. Osawa, H. Matsubara, I. Noda, and M. Asada. The robocup synthetic agent challenge 97. In *Proceedings of IJCAI-97*, Nagoya, Japan, August 1997.
- [19] Mary Koes, Illah Nourbakhsh, and Katia Sycara. Heterogeneous multirobot coordination with spatial and temporal constraints. In *Proceedings of the Twentieth National Conference on Artificial Intelligence (AAAI)*.

- [20] J.L. Krichmar and G.M. Edelman. Brain-based devices: Intelligent systems based on principles of the nervous system. In *In Proceedings of the Conference on Intelligent Robots and Systems (IROS)*, Las Vegas, NV, 2003.
- [21] Maja J. Mataric. *Interaction and Intelligent Behavior*. PhD thesis, EECS, MIT, Boston, MA, August 1994. Available as technical report AITR-1495.
- [22] H. G. Nguyen, J. Morrell, K. Mullens, A. Burmeister, S. Miles, K. Thomas, and D. W. Gage. Segway robotic mobility platform. In *SPIE Mobile Robots XVII*, October 2004.
- [23] Monica Nicolescu and Maja J Mataric. 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*, pages 241–248, 2003.
- [24] I. Nourbakhsh, E. Hamner, E. Porter, B. Dunlavy, E. Ayoob, T. Hsiu, M. Lotter, and S. Shelly. The design of a highly reliable robot for unmediated museum interaction. In *Proceedings of ICRA '05, the 2005 IEEE International Conference on Robotics and Automation*, 2005.
- [25] L. E. Parker, B. Kannan, F. Tang, , and M. Bailey. Tightly-coupled navigation assistance in heterogeneous multi-robot teams. In *In Proceedings of the Conference on Intelligent Robots and Systems (IROS)*, 2004.
- [26] Lynne E. Parker. Alliance: An architecture for fault tolerant multi-robot cooperation. *IEEE Transactions on Robotics and Automation*, 14 (2), 1998.
- [27] D. Polani, B. Browning, A. Bonarini, and K. Yoshida, editors. *RoboCup 2003: Robot Soccer World Cup VII*. Springer-Verlag Press, Berlin, 2003.
- [28] M. Roth, R. Simmons, and M. Veloso. Decentralized communication strategies for coordinated multi-agent policies. In *In Multi-Robot Systems: From Swarms to Intelligent Automata*, volume IV. Kluwer Academic Publishers, 2005.
- [29] S.Arulampalam, S.Maskell, N.Gordon, and T.Clapp. A tutorial on particle filters for on-line non-linear/non-gaussian bayesian tracking. *IEEE Transactions on Signal Processing*, 50(2):174–188, February 2002.

- [30] Paul Scerri, Y. Xu, Elizabeth Liao, J. Lai, and Katia Sycara. Scaling teamwork to very large teams. In *AAMAS'04*, 2004.
- [31] J. Searock, B. Browning, and M. Veloso. Turning segways into robust human-scale dynamically balanced soccer robots. July 2004.
- [32] Douglas Vail and Manuela Veloso. Dynamic multi-robot coordination. In *Multi-Robot Systems*. Kluwer, 2003.
- [33] M. Veloso, E. Pagello, and H. Kitano, editors. *RoboCup-99: Robot Soccer World Cup III*. Springer-Verlag Press, Berlin, August 2000.
- [34] Manuela Veloso and Peter Stone. Individual and collaborative behaviors in a team of homogeneous robotic soccer agents. In *Proceedings of the Third International Conference on Multi-Agent Systems*, pages 309–316, Paris, July 1998.
- [35] T.Kirubarajan Y.Bar-Shalom, X.-R.Li. *Estimation with Applications to Tracking and Navigation*. John Wiley & Sons, Inc, 2001.