

Coaching Robots to Play Soccer via Spoken-Language

Alfredo Weitzenfeld¹, Carlos Ramos¹, Peter Ford Dominey²

¹ ITAM, San Angel Tizapán, México DF, CP 0100
alfredo@itam.mx

<http://www.cannes.itam.mx/Alfredo/English/Alfredo.htm>

² Institut des Sciences Cognitives, CNRS, 67 Blvd. Pinel, 69675 Bron Cedex, France
dominey@isc.cnrs.fr

<http://www.isc.cnrs.fr/dom/dommenu-en.htm>

Abstract. The objective of this paper and our current research is to develop a human-robot interaction architecture that will let human coaches train robots to play soccer via spoken language. This work exploits recent developments in cognitive science, particularly notions of grammatical constructions as form-meaning mappings in language, and notions of shared intentions as distributed plans for interaction and collaboration between humans and robots linking perceptions to action responses. We define two sets of voice-driven commands for human-robot interaction. The first set involves action commands requiring robots to perform certain behaviors, while the second set involves interrogation commands requiring a response from the robot. We then define two training levels to teach robots new forms of soccer-related behaviors. The first level involves teaching new *basic* behaviors based on action and interrogation commands. The second level involves training new *complex* behaviors based on previously learnt behaviors. We explore the two coaching approaches using Sony AIBO robots in the context of RoboCup soccer standard platform league previously known as the four-legged league. We describe the coaching process, experiments, and results. We also discuss the state of advancement of this work.

1 Introduction

We expect interaction between humans and robots to be as natural as interaction among humans. As part of this work we are developing a domain independent language processing system that can be applied to arbitrary domains while having psychological validity based on knowledge from social cognitive science. In particular our architecture exploits models on: (i) language and meaning correspondence relevant to both neurological and behavioral aspects of human language developed by Dominey et al. [1], and (ii) perception and behavior correspondence based on the notion of shared intentions or plans developed by Tomasello et al. [2, 3]. In Weitzenfeld and Dominey [4] we describe preliminary results using two prior robotic platforms, the Event Perceiver System and AIBO robots under simpler tasks. The current paper reports more advanced yet preliminary results in coaching robots to play soccer in the context of RoboCup [5] standard platform league where ITAM's Eagle Knights team regularly competes at the top

level [6, 7]. In the standard platform league two teams of four fully autonomous robots in red or blue uniform play soccer on a 6m by 4m carpeted soccer field using Sony's AIBO robots (up until 2008). Robots have as main sensor a color-based camera with main actuators the four robot legs. The AIBO also includes wireless communication capabilities to interact with the game controller and other robots in the field. The field includes two colored goals and two colored cylinders used for localization. In order to win robots need to score as many goals as possible by kicking the orange ball into the opposite team's goal. As in human soccer, teams need to outperform the opponents by moving faster, processing images more efficiently, localizing and kicking the ball more precisely, and having more advanced individual and team behaviors. In general, approaches to robot programming vary from direct programming to learning approaches that are usually intended to optimize individual robot tasks. Our human-robot interaction approach is intended to coach robots to play soccer initially as individuals and then as a team as real human coaches currently do. The following sections describe the human-robot interaction architecture, the coaching approach, current experimental results, and conclusions.

2 Human-Robot Interaction Architecture

In this section we describe the human-robot interaction spoken-language architecture we have developed using the CSLU Speech Tools Rapid Application Development (RAD) system [8] for: (a) scene processing for event recognition, (b) sentence generation from scene description and response to questions, (c) speech recognition for posing questions, (d) speech synthesis for responding, and (e) sending and receiving textual communications with the robot.

2.1 Learning <Sentence, Meaning> Bindings

Dominey and Boucher [9, 10] describe a system that can adaptively acquire a limited grammar by training with human narrated video events. An image processing algorithm extracts the meaning of the narrated events translating them into action descriptors, detecting physical contacts between objects, and then using the temporal profile of contact sequences in order to categorize the events (see [11]). The visual scene processing system is similar to related event extraction systems that rely on the characterization of complex physical events (e.g. give, take, stack) in terms of composition of physical primitives such as contact (e.g. [12, 13]). Each narrated event generates a well formed <*sentence, meaning*> pair that is used as input to a model that learns the sentence-to-meaning mappings as a form of template where nouns and verbs can be replaced by new arguments in order to generate the corresponding new meanings. Each grammatical construction corresponds to a mapping from sentence to meaning. This information is also used to perform the inverse transformation from meaning to sentence. These templates or grammatical constructions (see [14]) are identified by the configuration of grammatical markers or function words within the sentences [15]. The construction set provides sufficient linguistic flexibility, for

example, when a human commands a robot to push a ball towards the goal *push(robot, ball)*.

2.2 Learning <Percept, Response> Bindings

Ad-hoc analysis of interaction among humans during teaching-learning reveals the existence of a general intentional plan that is shared between teachers and learners. In a generalized independent learning platform is described where new <percept, response> constructions can be acquired by binding together perceptual and behavioral capabilities. Three components are involved in these <percept, response> constructions: (i) the percept, either a verbal command or an internal system state, originating from vision or another sensor; (ii) the response to this percept, either a verbal response or a motor response from the existing behavioral repertoire; and (iii) the binding together of the <percept, response> construction and its subsequent validation that it was correctly learned. The system then links and saves the <percept, response> pair so that it can be used in the future.

Having human users control and interrogate robots through spoken language results in the ability to naturally teach robots individual action sequences conditional on perceptual values or even more sophisticated shared intention tasks involving multiple robots such as passing the ball between robots when one of them is blocked or far away from the goal. In Dominey and Weitzenfeld [15] we describe a general approach for action, interrogation and teaching robots simple tasks.

3 Coaching Robots to Play Soccer

In order to demonstrate the generalization of the human-robot interaction architecture we have begun a series of experiments in the domain of RoboCup Soccer a well documented and standardized robot environment that provides a quantitative domain for evaluation of success. While no human intervention is allowed during a game, in the future humans could play a decisive role analogous to real soccer coaches adjusting in real-time their team playing characteristics according to the state of the game, individual or group performance, or the playing style of the opponent. Furthermore, a software-based coach may become incorporated into the robot program analogous to the RoboCup simulated coaching league where coaching agents can learn during a game and then advice virtual soccer agents how to optimize their behavior accordingly (see [16, 17]).

3.1 Coaching System

The coaching system developed by our group is illustrated in Figure 1. The spoken language interface is provided by the CSLU-RAD while communication to the Sony AIBO robots is done in a wireless fashion via UDP messages. Messages communicate the EK2007 system with MsgToAIBO module in the remote computer. The UDP messages contain a data structure that can be used to send textual commands to the

robot or interrogate the robot about internal states. The *MsgToAIBO* module broadcasts the UDP messages containing information about commands and interrogations.

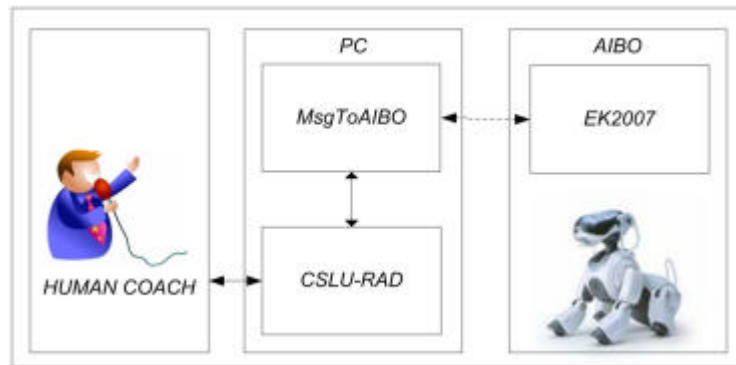


Fig. 1. Coaching system. A human coach interacts with the CSLU-RAD system to command or interrogate the AIBO robot. Textual messages are sent between the EK2007 system controlling the AIBO and the *MsgToAIBO* module in the remote computer.

The human coach uses voice commands to control the behavior of the AIBO. These voice commands are translated into regular text commands by the CSLU-RAD system and then sent transmitted to the EK2007 system controlling the behavior of the robot.

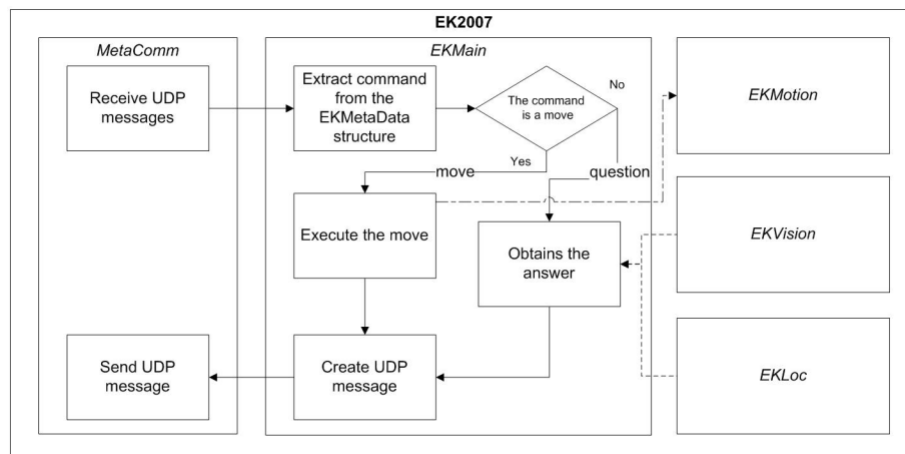


Fig. 2. EK2007 system. Commands are sent and receive via *MetaComm*. *EKMain* processes all commands and questions interacting with *EKMotion*, *EKVision* and *EKLoc* accordingly.

The EK2007 system is shown in more detail in Figure 2. It contains the following modules: *EKMain* controlling robot behavior; *EKMotion* controlling robot movement; *EKVision* processing camera images to identify objects of interest; *EKLoc* computing robot position in the field; and *MetaComm* sending and receiving UDP messages between robots and computer. The *MsgToAIBO* module takes as parameter a

MetaData containing information that includes an action or interrogation command and a destination robot identifier. When a UDP messages is received by the robot, the *MetaComm* module processes the *MetaData* verifying that the message is intended for the current robot, in which case it is delivered to the *EKMain* module. If the command refers to an action command, an action control is sent to *EKMotion*; otherwise if it refers to a question, information is requested from either *EKVision* or *EKLoc*. *EKMain* responds with a UDP message sent back to the computer via the *MetaComm* module. Action commands finish immediately while interrogation commands wait for an answer.

3.2 Action and Interrogation Commands

In order to demonstrate the human-robot coaching architecture we have defined and developed a number of basic action and interrogation commands in the context of soccer playing robots. These actions and interrogations are executed as voice commands using the CSLU-RAD software.

Table 1 describes action commands showing corresponding *voice command*, *CSLU node* (see Figure 3), *command id*, and *behavior*. Action commands include *Stop*, *Walk*, *Turn Right*, *Turn Left*, *Kick Ball*, *Hold Ball*, *Turn Right with Ball*, *Turn Left with Ball*, *Go to Ball*, and *Block Ball*. Note that certain actions such as *Go to Ball* depend on perceptions, in this case seeing the ball.

Table 1. Action commands showing voice command, corresponding CSLU node, command id and behavior.

<i>Voice Command</i>	<i>CSLU node</i>	<i>Command Id</i>	<i>Behavior</i>
Stop	Stop	0	Stop
Walk	Walk	1	Walk
Right	TurnR	9	Turn Right
Left	TurnL	10	Turn Left
Kick	Kick	11	Kick Ball
Hold Ball	Hold	12	Hold Ball
Right with Ball	TRH	13	Turn Right with Ball
Left with Ball	TLH	14	Turn Left with Ball
Go to Ball	Go -> Ball	20	Go to Ball and Stop in Front of Ball
Block	Block	22	Block Ball

Table 2. Interrogation commands showing voice command, corresponding CSLU node, command id and response.

<i>Voice Command</i>	<i>CSLU node</i>	<i>Command Id</i>	<i>Response - Id</i>
Do you see the ball?	Ball?	50	sees the ball – 1, otherwise - 0
Do you see the yellow goal?	YGoal?	51	sees the yellow goal – 1, otherwise - 0
Do you see the cyan goal?	CGoal?	52	sees the cyan goal – 1, otherwise - 0
Is the ball near?	BallN?	59	sees the ball near – 1, otherwise - 0

Table 2 describes the interrogation commands showing corresponding *voice command*, *CSLU node* (see Figure 3), *command id*, and *response*. Interrogation commands include *Do you see the ball?*, *Do you see the yellow goal?*, *Do you see the cyan goal?*, and *Is the ball near?*

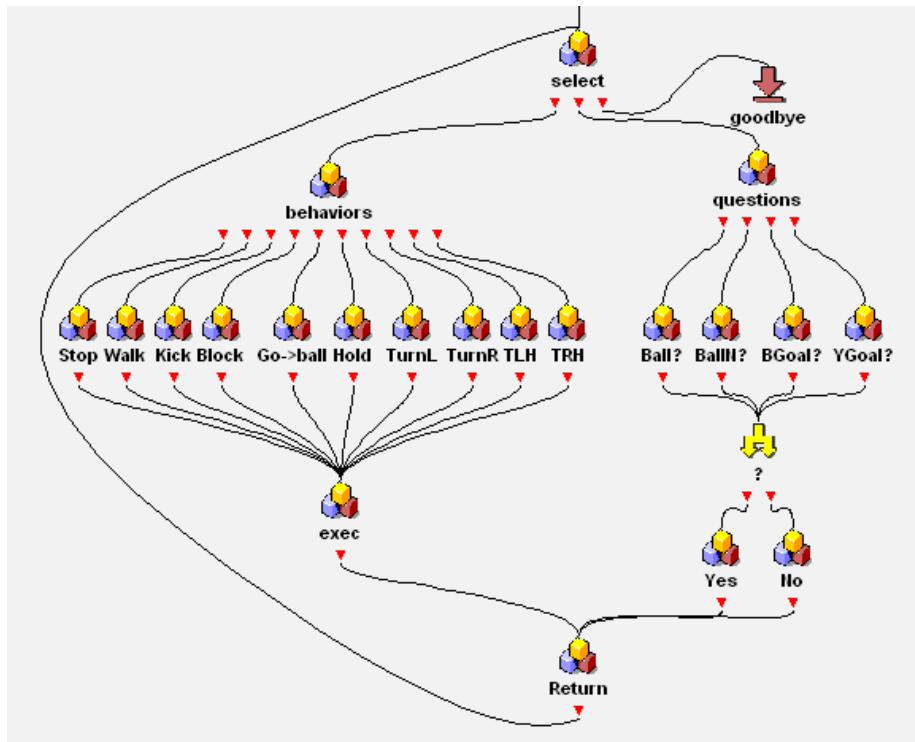


Fig. 3. The CSLU-RAD diagram describes the basic set of behaviors and questions that can be sent as voice commands to the robot.

CSLU-RAD defines a directed graph where each node in the graph links voice commands to specific behaviors and questions sent to the robot as shown in Figure 3. The *select* node separates action and interrogation commands. Action commands are represented by the *behaviors* node ('behavior' voice command) while interrogation commands are represented by the *questions* node ('question' voice command). Behavior nodes include 'Stop', 'Walk', 'Kick', 'Go->ball', 'Hold', 'TurnL', 'TurnR', 'TLH', and TRH; while question nodes are 'Ball?' ('Do you see the ball?'), 'BallN?' ('Is the ball near?'), 'CGoal' ('Do you see the cyan goal?') and 'YGoal' ('Do you see the yellow goal?'). Behavior commands are processed by the *exec* node while questions are processed by the ? node that waits for a 'Yes' or 'No' reply from the robot. Finally, the *Return* node goes back to the top of the hierarchy corresponding to the *select* node. The *goodbye* node exits the system.

3.3 Teaching New Behaviors

Action and interrogation commands form the basis for teaching new behaviors in the system. In particular, we are interested in teaching soccer-related tasks at two levels: (i) basic behaviors linking interrogations to actions such as *if you see the ball then go to the ball*, or *if you are near the ball then shoot the ball*; and (ii) complex behaviors composed of previously learnt behaviors such as *“Go to the ball and shoot the ball”*.

The teaching process is based on individual action and interrogation commands. To achieve teaching, we have extended the CSLU-RAD model shown in Figure 3 to enable creation of new behavior sequences. These sequences are stored as files and then recalled as new behaviors. The updated voice command diagram is shown in Figure 4. The main difference with the previous diagram is that after the *questions* node the model saves the response (*Save* node) and continues directly to the *behaviors* node where actions are taken and then saved as part of the teaching process (*Save&Return* node). Additionally, all sequences learnt in the models are included as new behaviors in the system. The new behaviors currently described in the paper and shown in the diagram correspond to *GO*, *SHOOT*, and *GO&SHOOT* nodes.

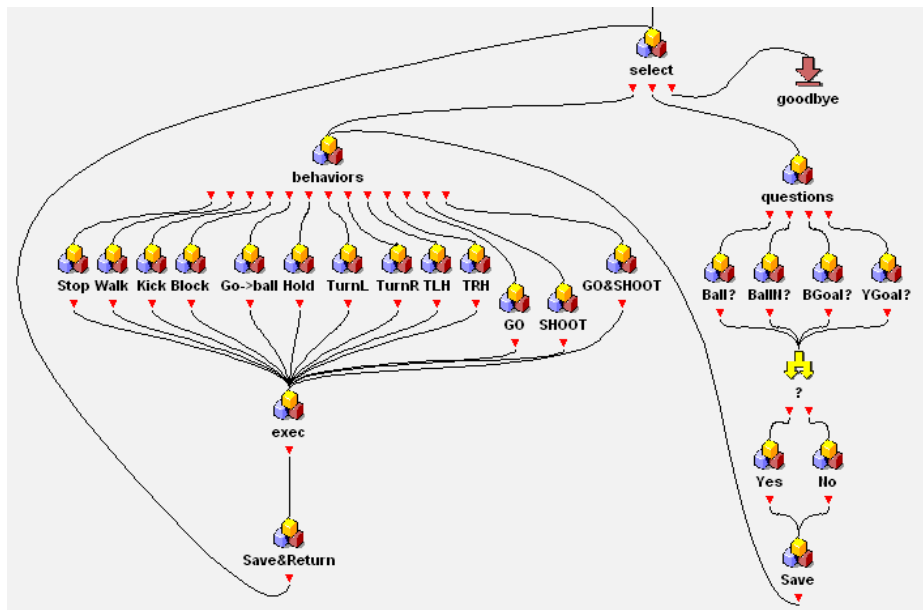


Fig. 4. The CSLU-RAD diagram describes the extended model for training the robot to Go and Shoot the ball.

A teaching conversation is represented by a sequence of action and interrogation commands sent to the robot. In Table 3 we show two teaching sequences: (i) *GO* training the robot to go towards the ball (left column); and (ii) *SHOOT* training the robot to kick the ball towards the goal (right column).

Table 3. Individual ‘Go’ (left column) and ‘Shoot’ (right column) training sequences.

‘GO’ training sequence	‘SHOOT’ training sequence
RAD: Select option User: Question User: Do you see the ball? RAD: No User: Turn right RAD: Select option User: Question User: Do you see the ball? RAD: Yes User: Go to the ball RAD: Select option User: Question User: Is the ball near? RAD: No User: Go to the ball RAD: Select option User: Question User: Is the ball near? RAD: Yes User: Hold the ball RAD: Select option User: Goodbye	RAD: Select option User: Question User: Do you see cyan goal? RAD: No User: Turn left RAD: Select option User: Question User: Do you see cyan goal? RAD: Yes User: Kick RAD: Select option User: Goodbye

‘GO’ begins with the robot not seeing the ball. The coach asks the robot to turn right until the ball is seen. The coach then asks the robot to walk towards the ball until the ball is near. When the ball is near, the coach asks the robot to hold the ball with its front legs and head. ‘SHOOT’ begins with the robot holding the ball. The coach asks the robot to keep turning right until it sees the cyan goal. The coach then asks the robot to kick the ball towards the goal. Note how we group interactions together into *interrogation-response-action* subsequences, e.g. (1) “User: Do you see the ball?”, (2) “RAD: No”, and (3) “User: Turn right”. These groups of three commands may be repeated several times as part of a single teaching conversation as in ‘GO’ and ‘SHOOT’ cases. New behaviors are then recalled as “User: Go” and “User: Shoot”.

There are two ways to build ‘GO&SHOOT’ training sequences: (i) use the *basic* sequence shown in Table 4 (left column) consisting of all the training commands used for ‘GO’ and ‘SHOOT’ added together; or (ii) use the *complex* sequence combining the learnt ‘GO’ and ‘SHOOT’ sequences as shown in Table 4 (right column). In the basic case, training proceeds as before by combining questions, responses and action commands. In the complex case, training involves simply recalling the already taught sequences i.e. ‘GO’ sequence followed by ‘SHOOT’.

Table 4. Two version of the combined ‘Go&Shoot’ training sequence: basic sequence (left) and complex sequence (right).

‘GO&SHOOT’ basic sequence	‘GO&SHOOT’ complex sequence
RAD: Select option User: Question User: Do you see the ball? RAD: No User: Turn right RAD: Select option User: Question User: Do you see the ball? RAD: Yes User: Go to the ball RAD: Select option User: Question User: Is the ball near? RAD: No User: Go to the ball RAD: Select option User: Question User: Is the ball near? RAD: Yes User: Hold the ball RAD: Select option User: Question User: Do you see cyan goal? RAD: No User: Turn left RAD: Select option User: Question User: Do you see cyan goal? RAD: Yes User: Kick RAD: Select option User: Goodbye	RAD: Select option User: Behavior User: Go RAD: Select option User: Behavior User: Shoot RAD: Select option User: Goodbye

4 Coaching Results

We tested the coaching architecture on the two basic behaviors described in the previous section: ‘GO’ and ‘SHOOT’, and then combined them into ‘GO&SHOOT’ in its basic and complex form. Table 5 shows the average user training time in seconds for the four sequences. Training time includes all voice interactions between human and robot as well as the time it takes the robot to perform the actual task. Note that ‘GO&SHOOT’ in its basic form is a direct sum of the individual ‘GO’ and ‘SHOOT’ training sequences. On the other hand, the average training time for ‘GO&SHOOT’ in its complex form is less than half since the task is taught through a more compact set of human-robot interactions based on the already learnt ‘GO’ and ‘SHOOT’ sequences. In other words, basic sequences take much longer to train than complex sequences since many more human-robot interactions are required to achieve the intended behavior, i.e. questions, response and action commands. The numbers shown in Table 5 are based on two to four trainings for basic behaviors and single trainings for complex behaviors.

Table 5. Average training time for ‘GO’, ‘SHOOT’, and ‘GO&SHOOT’ basic and complex sequences.

Average Training Time (sec)			
GO	SHOOT	GO&SHOOT - basic	GO&SHOOT - complex
63	29	93	40

Table 6. Test time for ‘GO&SHOOT’ behaviors with the robot and ball located at different initial positions. Both basic and complex ‘GO&SHOOT’ training sequences performed similarly .

Test	Robot Initial field position	Ball field position	Time (sec)				Total	Kick Result
			GO		SHOOT			
			looking for the ball	going to the ball	looking for the goal	kicking		
1	Left	Left	0	8	2	4	14	Goal
2	Center	Left	0	10	2	4	16	Goal
3	Right	Left	4	14	2	3	23	Fail
4	Left	Center	2	10	1	4	17	Fail
5	Center	Center	0	10	0	4	14	Goal
6	Right	Center	7	14	0	3	24	Goal
7	Left	Right	2	15	4	3	24	Goal
8	Center	Right	6	11	5	4	26	Goal
9	Right	Right	0	10	5	3	18	Goal

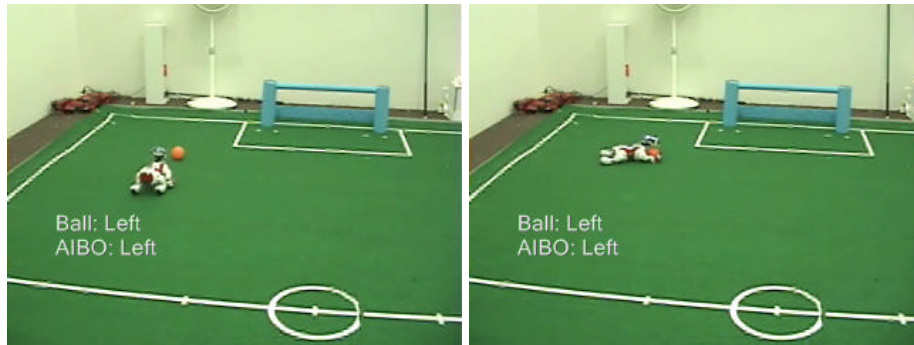


Fig. 5. The images show the behavior of the AIBO robot during ‘GO&SHOOT’ testing task. The left image shows part of the ‘GO’ task while the right image shows part of the ‘SHOOT’ task. The task was tested with ball and robot at different initial positions.

Table 6 shows the result from testing the ‘GO&SHOOT’ task in the standard platform soccer field under real lighting conditions. Both basic and complex forms performed comparably. The test involves going for the ball and shooting it into the goal with different initial positions for both ball and robot. We test on half of the field. The part of the task that takes the longest is going to the ball. Note how time varies depending on the robot and ball relative initial position, i.e. takes more time to complete the task if robot and ball are on different sides of the goal. Additionally, tests take longer if the robot does not see the ball immediately. We also include in the table the kick results, i.e. goal or fail. Note that scoring the goal depends on the robot correct orientation and kicking.

Figure 5 shows two images taken from ‘GO&SHOOT’ testing tasks. The figure in the left shows the robot going to the ball, part of ‘GO’ task, while the figure in the right shows the robot shooting the ball, part of the ‘SHOOT’ task. The full video for training and testing sets can be found in [18].

Conclusions and Discussion

We have described in this paper results from our current research in the development of a generalized approach to human-machine interaction via spoken language in the context of robot soccer. The coaching architecture previously described exploits recent developments in cognitive science - particularly notions of grammatical constructions as form-meaning mappings in language, and notions of shared intentions as distributed plans for interaction and collaboration binding perceptions to actions. With respect to social cognition, shared intentions represent distributed plans in which two or more collaborators have a common representation of an action plan in which each plays specific roles with specific responsibilities with the aim of achieving some common goal. In the current study, the common goals were well defined in advance (e.g. teaching the robots new relations or new behaviors), and so the shared intentions could be built into the dialog management system.

To demonstrate the interaction model we described in the current work how to coach a robot to play soccer by teaching new behaviors at two levels: (i) basic behaviors trained from a sequence of existing actions and interrogations, and (ii) complex behaviors trained from newly trained sequences. We tested the architecture on ‘GO’ and ‘SHOOT’ tasks using Sony AIBOs in the context of RoboCup soccer standard platform league where our ITAM Eagle Knights team consistently competes at the top level. This work has demonstrated our technical ability to teach robots new behaviors, yet we need to provide with a more natural interaction between humans and robots. In particular the dialog pathways are somewhat constrained, with several levels of hierarchical structure in which the user has to navigate the control structure with several single word commands in order to teach the robot a new relation, and then to demonstrate the knowledge, rather than being able to do these operations in more natural single sentences. In order to address this issue, we are reorganizing the dialog management where context changes are made in a single step. Also, in order to focus the interactions, we are working around scenarios in which the human can interact with several robots at once collaborating around a shared goal such as passing the ball between themselves as part of more sophisticated game playing. In such scenario the human coach will be eventually be able to transmit knowledge in the form “*if blocked pass the ball to player behind*”. Such a command will modify an internal robot database with “*if possess(ball) and goal(blocked) then pass(ball)*”.

Our long term goal in human-robot coaching is to be able to positively affect team performance during a real game. Eventually, these coaching capabilities will be done by an agent coaches inside the robot.

Acknowledgements

Supported by French-Mexican LAFMI, ACI TTT Projects in France, UC-MEXUS CONACYT, CONACYT #42440, and “Asociación Mexicana de Cultura” in Mexico.

References

1. Dominey PF, Hoen M, Lelekov T and Blanc JM, Neurological basis of language in sequential cognition: Evidence from simulation, aphasia and ERP studies, *Brain and Language*, 86(2):207-25, 2003.
2. Tomasello M, *Constructing a language: A usage-based theory of language acquisition*. Harvard University Press, Cambridge, 2003.
3. Tomasello M, Carpenter M, Call J, Behne T, Moll H, *Understanding and sharing intentions: The origins of cultural cognition*, Behavioral and Brain Sciences, 2006.
4. Weitzenfeld, A., and Dominey, P., 2007, Cognitive Robotics: Command, Interrogation and Teaching in Robot Coaching, *RoboCup 2006: Robot Soccer World Cup X*, G. Lakemeyer et al. (Eds.), LNCS 4434, pp. 379–386, Springer-Verlag
5. Kitano H, Asada M, Kuniyoshi Y, Noda I, and Osawa, E.. Robocup: The robot world cup initiative. In *Proceedings of IJCAI-95 Workshop on Entertainment and AI/ALife*, 1995.
6. Weitzenfeld, A., Martínez, A., Muciño, B., Serrano, G., Ramos, C., and Rivera C., 2007, *EagleKnights 2007: Four-Legged League*, Team Description Paper, ITAM, Mexico.
7. Martínez-Gómez J.A. and Weitzenfeld A, Real Time Localization in Four Legged RoboCup Soccer, *Proc. 2nd IEEE-RAS Latin American Robotics Symposium*, Sao Luis Maranhao Brasil, Sept 24-25, 2005.
8. CSLU Speech Tools Rapid application Development (RAD), <http://cslu.cse.ogi.edu/toolkit/index.html>
9. Dominey PF and Boucher JD, Developmental stages of perception and language acquisition in a perceptually grounded robot, *Cognitive Systems Research*, Volume 6, Issue 3, Pages 243-259, September 2005.
10. Dominey PF and Boucher JD, Learning to talk about events from narrated video in a construction grammar framework, *AI*, Vol 167, No 1-2, pp 31-61, Sept 2005.
11. Kotovsky L and Baillargeon R, The development of calibration-based reasoning about collision events in young infants, *Cognition*, 67, 311-351, 1998.
12. Siskind JM, Grounding the lexical semantics of verbs in visual perception using force dynamics and event logic. *Journal of AI Research* (15) 31-90, 2001.
13. Steels L and Baillie JC. Shared Grounding of Event Descriptions by Autonomous Robots. *Robotics and Autonomous Systems*, 43(2-3):163—173, 2002.
14. Goldberg A, *Constructions*. U Chicago Press, Chicago and London, 1995.
15. Bates E, McNew S, MacWhinney B, Devescovi A, and Smith S, Functional constraints on sentence processing: A cross linguistic study, *Cognition* (11): 245-299, 1982.
16. Riley P, Veloso M, and Kaminka G, An empirical study of coaching. In: *Distributed Autonomous Robotic Systems 6*, Spring-Verlag, 2002.
17. Kaminka G, Fidanboyly M, Veloso M, Learning the Sequential Coordinated Behavior of Teams from Observations. In: *RoboCup-2002 Symposium*, Fukuoka, Japan, June, 2002.
18. Videos coaching AIBO robots to play soccer, <http://robotica.itam.mx/EKCoaching/>