

## Instruction Based Learning: how to instruct a personal robot to find HAL.

Stanislao Lauria, Guido Bugmann<sup>1</sup>, Theodoros Kyriacou, Ewan Klein\*

Centre for Neural and Adaptive Systems, School of Computing, University of Plymouth  
Drake Circus, Plymouth PL4 8AA, United Kingdom.

\*Institute for Communicating and Collaborative Systems, Division of Informatics, University of Edinburgh, 2  
Buccleuch Place, Edinburgh EH8 9LW, Scotland, United Kingdom.

<http://www.tech.plym.ac.uk/soc/staff/guidbugm/ibl/index.html>

### Abstract

Future domestic robots will need to adapt to the special needs of their users and to their environment. Programming by natural language will be a key method enabling computer language-naïve users to instruct their robots. Its main advantages over other learning methods are speed of acquisition and ability to build high level symbolic rules into the robot. This paper describes the design of a practical system that uses unconstrained speech to teach a vision-based robot how to navigate in a miniature town. The robot knows a set of primitive navigation procedures that the user can refer to when giving route instructions.

Since the user is likely to refer to a procedure that is not pre-programmed in the robot, the system must be able to learn it. This paper investigates how to make the learning process possible. In particular, a method is proposed to fasten the choice of an initial set of primitives to the natural human speech chunking. Moreover, the use of Instruction-Based Learning (IBL) imposes a number of constraints on the design of robotics systems and knowledge representation. These issues are developed in the paper and proposed solutions described.

### 1. Introduction

Intelligent robots must be capable of action in reasonably complicated domains with some degree of autonomy. This requires adaptivity to a dynamic environment, ability to plan and also speed in the execution. In the case of helper robots, or domestic robots, the ability to adapt to the special needs of their users is crucial. The problem addressed here is the one of how a user could instruct the robot to perform tasks which manufacturers cannot

completely program in advance. In this case the system will not work at all if it cannot learn.

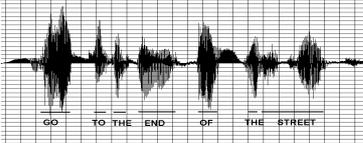
Such learning requires interaction and collaboration between the user and the robot. But, as most users are computer-language-naïve, they cannot personalise their robot using standard programming methods. Indirect methods, such as learning by reinforcement or learning by imitation, are also not appropriate for acquiring user-specific knowledge. For instance, learning by reinforcement is a lengthy process that is best used for refining low-level motor controls, but becomes impractical for complex tasks. Further, both methods do not readily generate knowledge representations that the user can interrogate.

Instruction-Based Learning (IBL), which uses unconstrained speech, has several potential advantages. Natural language can express rules and sequences of commands in a very concise way. Natural language uses symbols and syntactic rules and is well suited to interact with robot knowledge represented at the symbolic level. It has been shown that learning in robots is much more effective if it operates at the symbolic level (Cangelosi and Harnad, 2001). This is to be contrasted with the much slower learning at the level of direct sensory-motor associations.

Chunking, sequencing and repair are the aspects, related to natural language interactions, shaping the design of IBL systems discussed here. Chunking is a principle that applies to the communication of information. Chunking is meant here as the human characteristic to divide, during explanations, tasks into sub-tasks, so that all information should be presented in small 'basic' units of actions. As shown in (Miller 1956), chunking is done spontaneously by humans and consequently the system must be on the same 'wavelength' as the user in order to be successful. This means establishing for the robot the appropriate

---

<sup>1</sup> To whom correspondence should be addressed.

Analysis			Repair
	Speech recognition		
↓	Tagging	Go/VB to/TO the/DT end/NN of/IN the/DT street/NN	↑
	Syntactic Parsing	[VG go] [to to] [NG the end of the street ]	
↓	Semantic Analysis	Robot ( x ), end_of_street ( y ), request_go ( x, y )	↑
↓	Functional Mapping	Goto("end_of_street")	↑
↓	Robot program	Until found(end_of_street) follow_the_road()	↑

**Table 1.** From speech to action. The various steps involved in the transformation of a user command into the corresponding action are shown here.

prerequisites for the conversion of cognition, carried in chunks, into the form of procedures. In a robot involved with navigation tasks, a fundamental prerequisite is that the system must possess a set of pre-programmed procedures related to the very basic chunks used in route instruction situations. Moreover, since in the learning process the user does not express his requirement with a single chunk, the system must be able to sequence the chunks correctly. For example, in a sequence of instructions given by the user, the final state of an action may not correspond to the expected state for the next action. In this case, the system would not be able to perform its task due to the missing chunk. For this reason, it is necessary to define a proper internal knowledge representation allowing the system to detect the missing information. In this way, the system would be able to make predictions about future events so that the problem can be solved while the system is still interacting with the user.

Finally, the system not only has to pay attention to user knowledge and dialogue goals, but it also has to adapt its dialogue behaviour to current limitations of the user's cognitive processing capabilities. Assistance is then expected from the system, so that the interaction may naturally flow over the course of several dialogue turns. Moreover, a dialogue manager should take care of identifying, and recovering from, speech recognition and understanding errors.

This paper describes initial steps and considerations towards a practical realisation of an IBL system. The experimental environment is that of a miniature town in which a robot provided with video camera executes route instructions. The robot has a set of pre-programmed sensory-motor action primitives, such as "turn left" or "follow the road". The task of the user is to teach the robot new routes

by combining action primitives. That task should reveal all the constraints described above, and enable testing of the developed methodology.

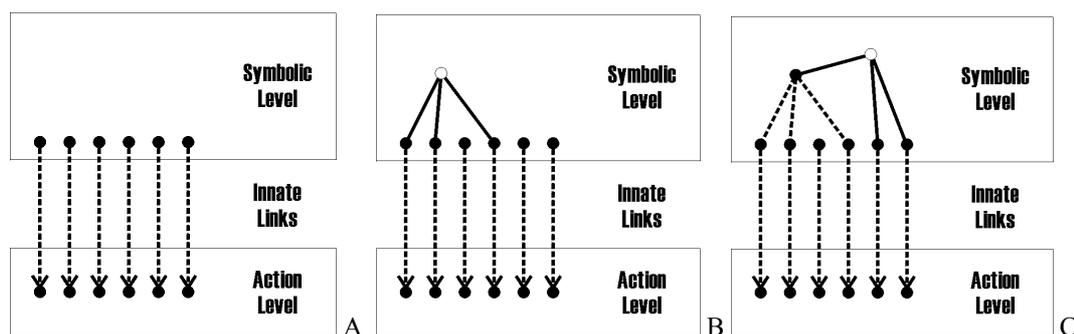
In the next section the IBL architecture implications due to chunking, sequencing and repair are discussed and how the rest of this paper is organized is also specified.

## 2. The big picture: from verbal utterance to robot action

With IBL, the system must convert verbal instructions given by the user into procedures containing internal program code controlling the robot sensors/actuators. It is during the learning process that such procedures are created and become part of a pool of procedures that can then be reused to learn more and more complex procedures. In this way the robot becomes able to execute increasingly complex tasks based on a set of pre-programmed primitives.

The closer the correspondence between primitives and chunks expressing the very basic actions (such as "turn left") is, the less difficult the learning is, since, in this way, the interaction between the user and system is kept to the minimum. For this reason, it is necessary to select these primitives that corresponds as closely as possible to the action expressed in the chunks.

Then, there is the problem of handling the chunks. In table 1, an example is given showing the various steps necessary to transform a user chunk into a robot action. First, the robot must be able to perform some speech recognition tasks in order to convert speech into text. After that, some syntactic parsing and semantic analysis is carried out. Then at the functional mapping level, the system must be able to transform the user utterance into internal



**Figure 1.** Symbolic learning. (A) is a schematic representation of the initial system, comprising symbols associated with pre-programmed (innate) primitive action procedures. In (B) the user has defined a new procedure (open circle) as a combination of symbols. The new symbol is grounded because it is a construct of grounded symbols. In (C), the user has defined a new procedure that combines a procedure previously defined by himself with primitive action procedures.

symbols that the robot can understand. By understanding we mean here that there is a correspondence between symbols and actions or real-world objects. In this way, the appropriate procedure can be called to act on the sensors/motors accordingly to the user intentions.

This multi-step approach has system-wide repercussions on the design of a robot control system. For example, the robot must be able to distinguish a command to be executed immediately from an instruction to be memorized. This requires context resolution at the natural language processing level. Moreover, the robot must be able to verify that the instruction can be converted into an executable procedure. It requires an internal representation of consequences of actions and the ability to verify the correct action sequencing. The robot must also be able to execute a command while listening to the user, and must cope with interruptions and inappropriate answers to its requests. This requires carefully designed system architecture. Some of the aspects discussed here are presented in more detail in the next sections. In particular, section 3 clarifies how symbol-level description and low-level sensory motor action procedures are integrated. The proposed representation of procedural knowledge is also described. In section 4 the system architecture is described.

The problems of considering the appropriate selection of action primitives is described in section 5 by analyzing recorded route instructions, and establishing a list of actions that are natural to users. The results of this investigation are also discussed. One of them is that the list of primitives may not be a closed one. The implications of that and other findings is discussed in section 6, along with the question of how the proposed system compares to other approaches. The conclusion follows in section 7.

### 3. IBL model

#### 3.1 Symbolic learning

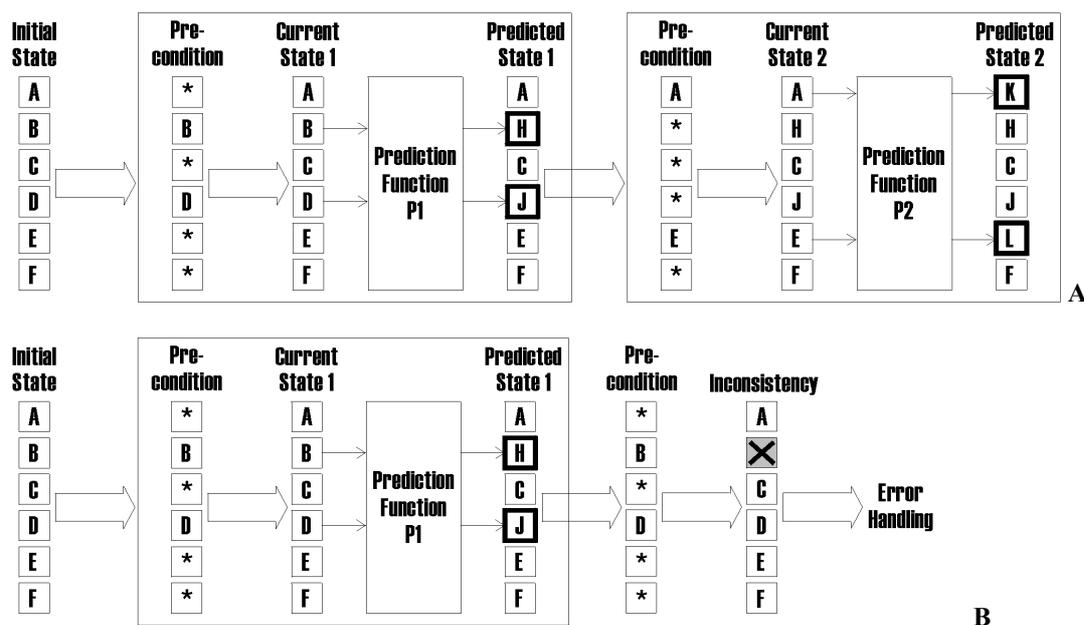
The learning process is based on predefined initial knowledge. This “innate” knowledge consists of primitive sensori-motor procedures with names, such as “turn left”, “follow the road” (the choice of these primitives is explained in sections 3.3 and 5). The name is what we call here a “symbol”, and the piece of computer program that controls the execution of the corresponding procedure is called the “action” (Figure 1A). As each symbol is associated with an action, it is said to be “grounded”.

When a user explains a new procedure to the robot, say a route from A to B that involves a number of primitive actions, the IBL system, on the one hand, creates a new name for the procedure, and, on the other hand, writes a new piece of program code that executes that procedure and links the code with the name (see section 3.2 for details). The code refers to primitive actions by name. It does not duplicate the low-level code defining these primitives. For that reason, the new program can be seen as a combination of symbols rather than a combination of actions (figure 1B). As all new procedures are constructed from grounded primitives, they become also grounded by inheritance and are “understandable” by the system when referred to in natural language.

When explaining a new procedure, the user can also refer to old procedures previously defined by himself. In that way the complexity of the robot's symbolic knowledge increases (fig. 1C).

#### 3.2 Knowledge representation

The internal representation needs to support three functions: (i) formal modeling of NL route descriptions; (ii) internal route planning for determining whether a given route description is sufficiently specified; and (iii) the generation of



**Figure 2.** Route instruction verification. (A) For each procedure there is a prediction function that transforms a state vector into its future value. The function first determines if the input state satisfied the minimal criteria ("pre-condition") to enable the procedure to be executed. An action is executable only if selected elements of the state vector have required values. If this is the case, the next state is predicted and processed by the prediction function associated with the next procedure in the instruction. Each action modifies certain components of the state vector, and leaves the other unchanged. (B) If the predicted state produced by one procedure does not allow the next procedure to be executed, an error handling process is initiated. (Note: the "initial state" in the text corresponds to the "current state" in the figure).

procedures for navigation at execution time. These three functions require different representations that will be described in turn.

(i) The utterances of the user are represented using the discourse representation structure (DRS) (Bugmann2001). This is then translated into symbols representing procedures or is used to initiate internal functions such as execution of a command or learning of a series of commands (section 4).

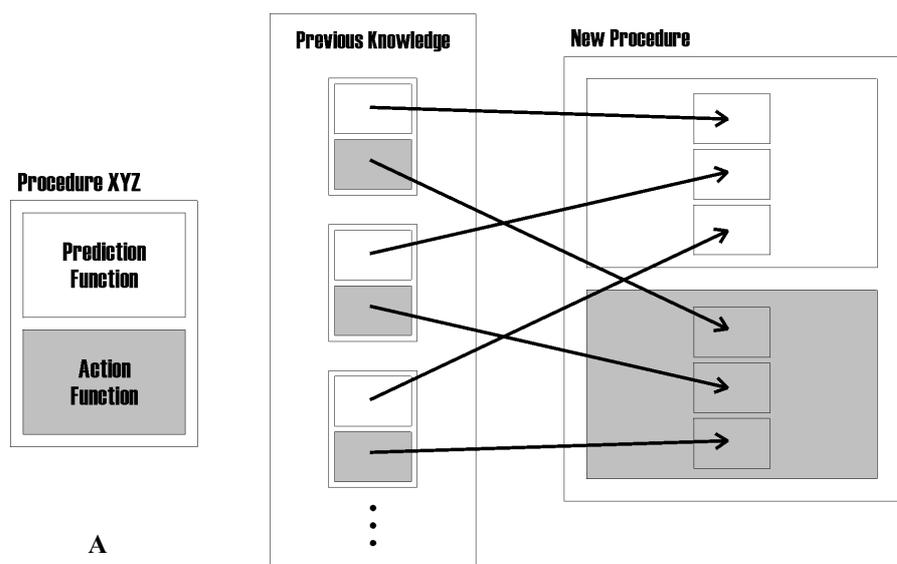
(ii) When the user describes a route as a sequence of actions, it is important for the robot to verify if this sequence is executable. The approach proposed here associate each procedure with a triplet  $S_i A_{ij} S_j$  with properties similar to productions in SOAR (Laird et al, 1987). The state  $S_i$  is the initial state in which the action  $A_{ij}$  can take place. It is the pre-condition for action  $A_{ij}$ . The state  $S_j$  is the final state, resulting of the action of  $A_{ij}$  applied to the initial state (figure 2 clarifies the difference between "initial state" and "pre-condition"). For a sequence of actions to be realisable, the final state of one action must be compatible with the pre-condition of the next one. To enable this verification, the robot must be able to "imagine" the consequence of an action. For that purpose, a PREDICTION function is associated with each primitive action, and with each newly created procedure. Figure 2 illustrates the use of the prediction function during verification of the

consistency of the sequence of instructions from the user. It should be noted that this process also helps detecting some of the errors in natural language processing.

(iii) When a robot executes a command, it executes a piece of program code that contains the sequence of primitive procedures to be executed. Thus, a key part of IBL is the generation of a program code. This is enabled by the use of a scripting language (section 4). This program is called the ACTION function. Both ACTION and PREDICTION functions are physically located in the same file that contains all information specific to a procedure. This is schematised in figure 3.

### 3.3 Sensory-Motor primitives

Sensory-motor primitives are defined as actions that users usually refer to in unconstrained speech (chunking). These are not low-level robot control actions, and often involve complex processing and planning. A task such as "approach that building at the end of the street" is a typical action that users ask the robot to do at the end of a route instruction, when the goal is in sight (section 5). It is a complex action involving visual detection of a building and of its entrance, its localisation in relation to the street, and planning of a route along the street. All this is easy for a human, but in many ways stretches the limits of robot "intelligence".



**Figure 3.** Procedural knowledge representation. **(A)** A procedure file contains an *ACTION* function that causes the physical displacement of the robot, and a *PREDICTION* function that calculates the future state of the robot resulting from the action. The *ACTION* is used during execution of a command, and the *PREDICTION* is used for consistency checking during the learning process. **(B)** An instruction by the user results in a “New Procedure” file being written. In this file, the actions components of the requested primitive procedures are combined (in the form of function calls) to create the new *ACTION* function, and the prediction components are combined to create the new *PREDICTION* function. This includes an additional procedure-specific pre-condition.

We see here that, by setting the boundaries between the symbolic level and the action level to be the same as the one found in natural language, the symbolic level processing has been simplified, but at the cost of an increased complexity of “low-level” procedures. These give the robot some autonomy in the execution of commands, as the execution details depend on the local conditions.

#### 4. System Architecture

The architecture is comprised of several functional processing modules (figure 4). These are divided into two major units: the Dialogue Manager (DM) and the Robot Manager (RM).

The DM and the RM are designed as two different processes based on asynchronous communication protocols. These processes run concurrently on different processors. In this way, the system can handle, at the same time, both the dialogue aspects of an incoming request from the user (i.e. speech recognition and semantic analysis, or detection of a “stop” command) and the execution of a previous user request (i.e. check if the request is in the system knowledge domain, and execute vision-based navigation procedures).

Two aspects are essential with this concurrent-processes approach. Firstly, to define an appropriate communication protocol between the two processes. Secondly, to define an appropriate

architecture for the RM and DM allowing the two processes to both communicate with each other while performing other tasks. At present a communication protocol based on sockets and context-tagged messages is evaluated.

Moreover, the system must also dynamically adapt itself to new user requests or to new internal changes, by being able to temporarily suspend or permanently interrupt some previous activity. For example the user may want to prevent the robot crashing against a wall and must therefore be able to stop the robot while the robot is driving towards the wall. Hence, the importance of a concurrent approach where the system constantly listens to the user while performing other tasks and at the same time being able to adjust the task if necessary.

The Dialogue Manager is a bi-directional interface between the Robot Manager and the user, either converting speech input into a semantic representation, or converting requests from the Robot Manager into dialogues with the user. Its components are run as different processes communicating with each other via a blackboard architecture. The RM must concurrently listen/send requests from/to the DM and try to execute them. For this reason a multi-threads approach has been used. Its communication interface is a process that only launches a message evaluation thread “Execution Process” and resumes listening to the DM. The execution process then starts an appropriate thread for executing a command, or

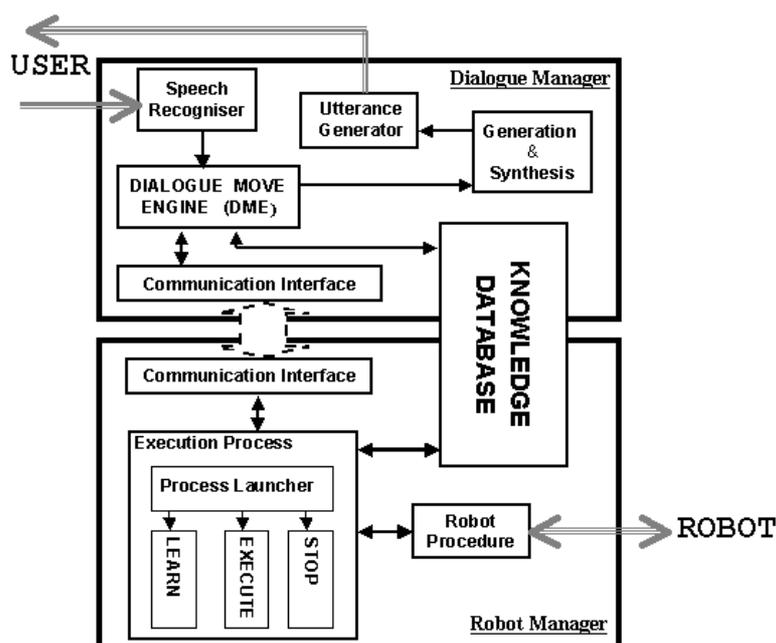


Figure 4. IBL system's architecture (see text for description).

places a tagged message on a message board if it is a part of a dialogue in a specific thread, e.g. learning a route. The characteristic of this approach is that all processes in the RM are sharing a common memory stack so that threads can be started and paused, depending on the user's input. At the moment, the Execution Process component is implemented with the Process Launcher controlling only the Learning and Execution modules since the Stop component is in an early stage of development. The Robot Manager is written using the scripting Python<sup>2</sup> language and C. An important feature of scripting languages is their ability to write their own code. For instance, a route instruction given by the user will be saved by the Robot Manager as a Python script that then becomes part of the procedure set available to the robot for execution or for future learning.

## 5. Corpus Collection and Data Analysis

To evaluate the potential and limitations of IBL, a real-world instructions task is used, that is simple enough to be realisable, and generic enough to warrant conclusions that hold also for other task domains. A simple route scenario has been selected, using real speech input and a robot using vision to execute the instructed route (see 5.1 below for more details). The first task in the project is to define the innate actions and symbols in the route instruction domain. For this reason, a corpus

of route descriptions has been collected from students and staff at the University of Plymouth. In section 5.2 and 5.3 corpus collection and data analysis are presented.

### 5.1 Experimental Environment

The environment is a miniature town covering an area of size 170cm x 120cm (figure 5). The robot is a modified RobotFootball robot<sup>3</sup> with an 8cm x 8cm base (figure 6A). The robot carries a CCD colour TV camera<sup>4</sup> (628 (H) x 582 (V) pixels) and a TV VHF transmitter. Images are processed by a PC that acquires them via with a TV capture card<sup>5</sup> (an example of such image is shown in figure 6B). The PC then sends motion commands by FM radio to the robot. During corpus collection, the PC is also used to record instructions given by subjects.

The advantage of a miniature environment is the ability to build a complex route structure in the limited space of a laboratory. The design is as realistic as possible, to enable subjects to use expressions natural for the outdoor real-size environment. Buildings have signs taken from real life to indicate given shops or utilities such as the post-office. However, the environment lacks some elements such as traffic lights that may normally be used in route instructions. Hence the collected corpus is likely to be more restricted than for outdoor route instructions. The advantage of using

<sup>3</sup> Provided by Merlin Systems  
(<http://www.merlinsystems.com/>)

<sup>4</sup> Provided by Allthings Sales and Services  
(<http://www.allthings.com.au/>)

<sup>5</sup> TV Card: Hauppauge WinTV GO

<sup>2</sup> <http://www.python.org>



**Figure 5.** Miniature town in which a robot will navigate according to route instructions given by users. Letters indicate the destinations and origins of various routes used in the experiment.

a robot with a remote-brain architecture (Inaba et al., 2000) is that the robot does not require huge on-board computing and hence can be small, fitting the dimensions of the environment.

### 5.2 Collection of a corpus of route instructions

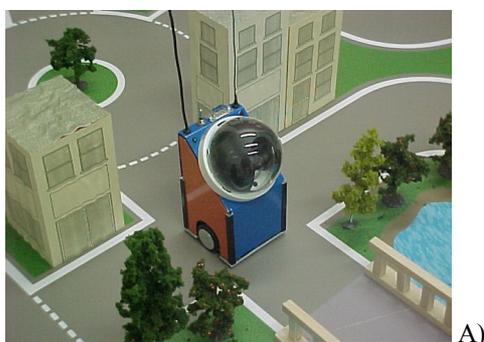
To collect linguistic and functional data specific to route learning, 24 subjects were recorded as they gave route instructions to the robot in the environment. Subjects were divided into three groups of 8. The first two groups (A and B) used totally unconstrained speech, to provide a performance baseline. It is assumed that a robot that can understand these instructions as well as a human operator would represent the ideal standard. Subjects from group C were induced in producing shorter utterances by a remote operator “taking notes”.

The other two groups (A and B) were told that the robot was remote-controlled and that, at a later date, a human operator would use their instructions

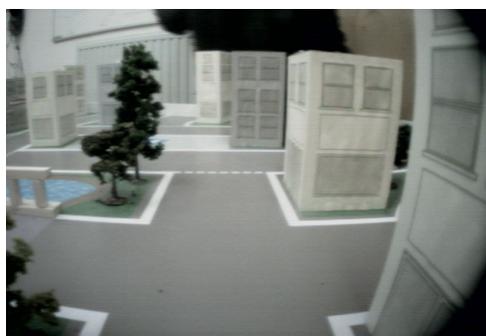
that the human operator would be located in another room, seeing only the image from the wireless on-board video camera. This induced subjects to use a camera-centred point of view relevant for robot procedure primitives. Subjects were also told to reuse previously defined routes whenever possible, instead of re-explaining them in detail. Each subject had 6 routes to describe among which 3 were “short” and 3 were “long”. The long routes included a short one, so that users could refer to the short one when describing the long one, instead of re-describing all segments of the short one. This was to reveal the type of expressions used by users to link taught procedures with primitive ones. Each subject described 6 routes having the same starting point and six different destinations. Starting points were changed after every two subjects. A total of 144 route descriptions were collected. For more details about collection and analysis of the corpus see (Bugmann et al. 2001).

### 5.3 Corpus Analysis: The functional vocabulary

The aim of the corpus analysis is to twofold. First, to define the vocabulary used by the users in this application, in order to tune the speech recognition system for an optimal performance in the task. Secondly, to establish a list of primitive procedures that users refer to in their instructions. The aim is to pre-program these procedures so that a direct translation from the natural language to grounded symbols can take place. In principle, if the robot does not know a procedure, the user could teach it. However, this is a process that we wish to avoid at this stage of the project, as discussed in section 6. Hereafter, we report on the functional analysis of the corpus. The reader interested in the task vocabulary can refer to (Bugmann et al., 2001). The functional vocabulary is a list of primitive navigation procedures found in route



A)



B)

**Figure 6** A. Miniature robot (base 8cm x 8cm). B. View from the on-board colour camera.

to drive the robot to its destination. It was specified

descriptions.

	Count	Primitive Procedures
1	308	MOVE FORWARD UNTIL [(past  over  across) <landmark>]   [(half_way_of   end_of) street]   [after <number><landmark> [left   right]]   [road_bend]
2	183	TAKE THE [<number>] turn [(left   right)]   [(before   after   at) <landmark>]
3	147	<landmark> IS LOCATED [left   right  ahead]   [(at   next_to   left_of   right_of   in_front_of   past   behind   on   opposite   near) < landmark >]   [(half_way_of   end_of   beginning_of   across) street]   [between <landmark> and <landmark>]   [on <number> turning (left   right)]
4	62	GO (before   after   to) <landmark>
5	49	GO ROUND ROUNDABOUT [left   right]   [(after   before   at) <landmark>]
6	42	TAKE THE <number> EXIT [(before   after   at) <landmark>]
7	12	FOLLOW KNOWN ROUTE TO <landmark> UNTIL (before   after   at) <landmark>
8	4	TAKE ROADBEND (left   right)
9	4	STATIONARY TURN [left   right   around]   [at   from <landmark>]
10	2	CROSS ROAD
11	2	TAKE THE ROAD in_front
12	2	GO ROUND <landmark> TO [front   back   left_side   right_side]
13	1	PARK AT <location>
14	1	EXIT [car_park   park]

**Table 2.** Primitive navigation procedures found in the route descriptions collected from groups A and C. Procedure 3 is used by most subjects to indicate the last leg of a route, when the goal is in sight.

The initial annotation of instructions in terms or procedures, as reported here, is somehow subjective, and influenced by two considerations. (i) The defined primitives will eventually be produced as C and Python Programs. It was hoped that only a few generic procedures would have to be written. Therefore, the corpus has been transcribed into rather general procedures characterised by several parameters (table 2). (ii) An important issue is knowledge representation. According to the SAS representation discussed in section 3.2, the executability of primitives can only be evaluated if their initial and final states are defined. Subjects however rarely specified explicitly the starting point of an action and sometimes did not define the final state in the same utterance. Nevertheless, it was assumed that the system would be able to infer the missing information from the context. Therefore, procedures without initial or final state were considered to be complete, and were annotated as such. The specifications of primitive procedures are likely to evolve during the project.

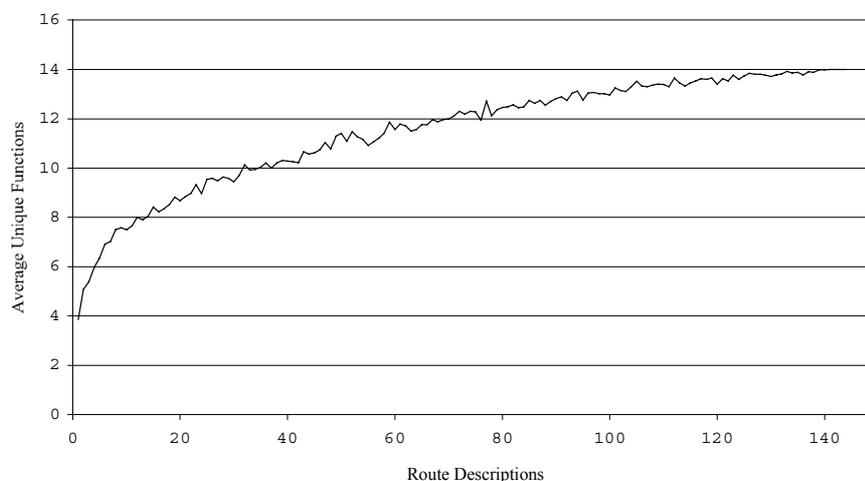
This analysis methodology differs slightly from the one in (Denis, 1997). In our analysis, there are no statements describing landmarks, as these are made part of procedures specifications, and consequently there are also no actions without reference to landmarks. Even when a subject specified a non-terminated action, such as "keep going", it was classified as "MOVE FORWARD UNTIL", assuming that a termination point would

be inferred from the next specified action. The list of actions found in the route descriptions of groups A and C is given in table 2. Figure 7 shows that the number of distinct procedures is increasing with the number of sampled instructions, but at a rate much smaller than the number of distinct words reported in (Bugmann et al., 2001). Here we discover on average one new procedure for every 38 route instructions, while with words, we discovered in average one new word for each instruction. New procedures typically are the least frequent in table 2.

## 6. Discussions

Teaching a route to a robot using natural language is an application of a more general instruction-based learning methodology. The corpus-based approach described here aims at providing users with the possibility of using unconstrained speech, whilst creating an efficient natural language processing system using a restricted lexicon. It is found that the functional vocabulary is small, containing only 12 primitives (although that number may vary with the annotation method). From a roboticist's point of view, route navigation could probably be achieved with a smaller number of primitives. However, when accepting spontaneous speech, a wider variety of functions must be expected.

An important finding is that the functional vocabulary is not closed. Hence, at some point in



**Figure 7.** Average number of unique procedures as a function of the number of collected route instructions. The curve is obtained by averaging 50 sets comprising a random selection of  $n$  route descriptions. The number  $n$  is shown on the  $x$ -axis of the graph. The slope of the curve indicates that, on average, one new function will be added to the functional lexicon for every 38 additional route instructions collected.

the robot's life, the user will have to teach it new primitives (e.g. "cross the road") or reformulate its instructions. To enable learning, the robot must possess a larger set of primitives, which correspond to lower level robot actions. For instance, the user may wish to refer to a number of wheel turns in its instruction. Examples of such instructions are found in (FLAKEY) and (Seabra Lopes, 2000). With our approach, this would require the collection of a new corpus to determine the necessary additional primitive procedures. Another solution may lie in an appropriate dialogue management to suggest a reformulation of the instruction. It is expected that with the corpus-based method used here, the frequency of such "repair dialogues" will be minimised. An open question is the detection of new functions in the user's utterance, as the lexicon may not contain the required vocabulary.

The approach to robot control described may be seen as an attempt to integrate the good properties of Behaviour-based control and classical AI. Behaviour-based control is an effective method for designing low-level primitives that can cope with real-world uncertainties, and AI has developed effective tools for symbol manipulation and reasoning (for a more detailed discussion about hybrid systems see for example Malcom (1995)). However, the system differs in several ways from both methods. Here, the corpus defines what symbols and primitives to use. Consequently, some of the primitives are rather complex functions, involving representations of the environment and planning. These are not always compatible with the representation-less philosophy of behaviour-based

systems. On the AI side, the system does not use the full range of reasoning capabilities offered by systems such as SOAR. There are no other aims in symbolic processing than verifying the consistency of instructions, and the construction of new procedure specifications.

Other previous work on verbal communication with robots has mainly focused on issuing commands, i.e. activating pre-programmed procedures using a limited vocabulary (e.g. IJCAI'95 office navigation contest). Only a few research groups have considered learning, i.e. the stable and reusable acquisition of new procedural knowledge. (Huffman & Laird, 1995) used textual input into a simulation of a manipulator with a discrete state and action space. (Crangle and Suppes, 1994) used voice input to teach displacements within a room and mathematical operations, but with no reusability. In (Torrance, 1995), textual input was used to build a graph representation of spatial knowledge. This system was brittle due to place recognition from odometric data and use of IR sensors for reactive motion control. Knowledge acquisition was concurrent with navigation, not prior to it. Whereas in (Matsui et al. 1999), the system could learn new actions through natural language dialogues but only while the robot was performing them (i.e. it could only learn a new route from A to B while it was actually moving from A to B and dialoguing with the user).

In the IBL system described here, learning operates purely at the symbolic level; hence it can be done prior to performance. The ability to predict future states enables to engage in a verification

dialogue before execution errors occur. If environmental conditions change such that an instruction is not valid anymore, this can be detected from the mismatch between the expected result and the actual one. Learning however is not autonomous. The system requires interaction with a human user to learn new symbols and their meaning. This simplifies the design of the robot due to the transfer of part of the cognitive load to the user. Future experiment will reveal if this approach results in effective and socially acceptable helper robots.

## 7. Conclusions

The project described in this paper is aimed at exploring IBL for route descriptions. It has been discussed how the design of the IBL system is adapted to natural human behaviour. Indeed, both the vocabulary matches to the unconstrained user language and the functional primitives built into the robot are determined from actions natural to the users. This defines an architecture open to spontaneous user interventions, unexpected replies and errors. Nevertheless, user-friendliness is not a prior specification, but a consequence of practical constraints. Indeed, robots without learning will not achieve specific tasks (such as finding HAL) and a system without adapted vocabulary causes too many errors. Similarly, explaining tasks is beyond the cognitive capabilities of users without high level primitives and, like with HAL, a robot that listens only when it decide to do so would be out of control. So far, the speech recognition part is in an early stage of development while the DRS part is operational for a limited number of examples, and that work is in progress to improve the coverage of corpus. However, it is found that the functional vocabulary is small, containing only 12 primitives (although that number may vary with the annotation method). The full transformation from NL utterances into procedures has been tested with dummy primitives (i.e. preprogrammed robot displacements). Programs for the proper sensory-motor primitives are currently under development. This will then allow further testing of the IBL concept.

However the initial results presented here show that neither the lexicon nor primitive procedures are likely to form closed sets. Ideally, IBL system should therefore also be capable of acquiring new words, and users should be given the possibility to teach new primitive 'innate' procedures. Unfortunately, the former is beyond the capabilities of current speech recognition systems. As for learning new primitives procedures, this would require a new set of more primitive procedures to

be combined via user instructions. Whether it will be possible to explore this during the project is unclear. To allow IBL to operate despite these limitations, it is likely that dialogue management will play a crucial role.

**Acknowledgement:** This work is supported by EPSRC grants GR/M90023 and GR/M90160. The authors are grateful to A. Cangelosi and K. Coventry for enlightening discussions.

## References:

- Bugmann G., Lauria S., Kyriacou T., Klein E., Bos J. and Coventry K. (2001) "Using Verbal Instruction for Route Learning", Proc. of 3rd British Conf. on Auton. Mobile Robots and Autonom. Systems: Towards Intelligent Mobile Robots (TIMR'2001), Manchester, 5 April.
- Cangelosi A., Harnad S. (2001) The adaptive advantage of symbolic theft over sensorimotor toil: Grounding language in perceptual categories. *Evolution Communication*. (in press)
- Crangle C. and Suppes P. (1994) Language and Learning for Robots, CSLI Lecture notes No. 41, Centre for the Study of Language and Communication, Stanford, CA.
- Denis M. (1997) "The description of routes: A cognitive approach to the production of spatial discourse", *CPC*, 16:4, pp.409-458.
- FLAKEY:  
[www.ai.sri.com/people/flakey/integration.html](http://www.ai.sri.com/people/flakey/integration.html)
- Huffman S.B. and Laird J.E. (1995) "Flexibly Instructable Agents", *Journal of Artificial Intelligence Research*, 3, pp. 271-324.
- Inaba M., Kagami S., Kanehiro F., Hoshino Y., Inoue H. (2000) "A platform for robotics research based on the remote-brained robot approach", *International Journal of Robotics Research*, 19:10, pp. 933-954.
- Laird J.E., Newell A. and Rosenbloom P.S. (1987) "Soar: An architecture for general Intelligence" *Artificial Intelligence*, 33:1, pp.1-64.
- Malcom C. M. (1995), The SOMASS system: a hybrid symbolic and behaviour-based system to plan and execute assemblies by robot. In J. Hallam, et al. (Eds), *Hybrid problems and Hybrid solutions* pp 157-168. Oxford: ISO-press.
- Matsui T., Asoh H., Fry J., et al. (1999) Integrated Natural Spoken Dialogue System of Jijo-2 Mobile Robot for Office Services, <http://citeseer.nj.nec.com/matsui99integrated.html>
- Miller G. (1956) "The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity Processing Information". *The Psychol. Review*, v. 63, p. 81-97
- Seabra Lopes, L. and A.J.S. Teixeira (2000) Human-Robot Interaction through Spoken Language Dialogue, Proceedings IEEE/RSJ International Conf. on Intelligent Robots and Systems, Japan.
- Torrance M.C. (1994) Natural Communication with Robots, MSc Thesis submitted to MIT Dept of Electrical Engin. and Comp. Science

