Anchoring symbols to sensorimotor *control*

Paul Vogt

IKAT - Universiteit Maastricht
P.O. Box 616, 6200 MD Maastricht, the Netherlands

Abstract

This paper investigates how robots may emerge a lexicon to communicate complex meanings about actions such as 'I am going to the red target' using simple (one-word) utterances. The main issue of the paper concerns the way these complex meanings represent the actions that are performed. It is argued that the meaning of these utterances may be represented without the need for categorising a complex flow of sensorimotor data. To illustrate the point, a simulation is presented in which robots develop such a communication system. The paper concludes by confirming that it is well possible to construct such a lexicon once robots have a number of basic sensorimotor skills available.

1 Introduction

In recent years an increasing number of studies have been done that investigate how a population of robots can develop a symbolic communication system. Although in most studies the communication system was given, e.g. [14], a few investigated how a communication system can emerge from scratch [10, 12, 13]. Each of these studies have had to tackle the *anchoring problem* [5] and the related *symbol grounding problem* [6]. Both problems relate to the question how symbols relate to the real world. The anchoring problem mainly relates to the technical issue of how to construct and maintain a relation (or *anchor*) between a symbol (or *word*) and the raw sensory data a robot senses concerning the word's reference. The symbol grounding problem is more related with fundamental and philosophical issues concerning the meaning of symbols.

A common approach for using symbols concerning actions is to anchor these on the flow of sensory data [5, 8, 12]. When using such techniques, many complications may occur because, for instance, the world changes continuously and sensors may be subject to much noise. This may cause many different interpretations of a particular scene or object when sensed at different occasions. A lot of diversity emerges in the categorisation that has to be disambiguated at the word-meaning level. The used models, such as the guessing game model, are capable to deal with disambiguating such diversity [13]. However, the diversity of interpretations may be reduced when symbols relate to the activation level of *control* mechanisms rather than to a flow of sensorimotor data.

In this paper, a model is presented for anchoring words to sensorimotor control, rather than to sensorimotor data. Especially when it comes to words that describe actions. This, then, provides another argument for the minimal representation of knowledge as advocated by the paradigm of behaviour-based robotics.

Taking up such an approach assumes the availability of such sensorimotor control modules (or *schemas* [1]). But when considering human language as an instance of symbolic communication, by the time human language evolved humans already had evolved many sensorimotor skills. This is what Brooks has meant when he noted that the development of functions such as using language would be "rather simple once the essence of being and reacting are available" [4].

2 The model

To test whether a symbolic communication system can emerge based on a minimal representation of the robots' actions, a simulation of a task-oriented experiment has been developed. The task the robots had to solve was to visit certain landmarks in their environment simultaneously. The robots had no notion of this task, but they would visit a landmark at certain moments in time and when they developed bits of a language, they tended to use the language to select their target. The robots developed a lexicon from scratch using the guessing game model, which is based on Steels' language game model and which has been implemented on real robots before [10, 12, 13].

The ideas were implemented in a mobile robot simulator in which an environment was designed containing four distinctively coloured landmarks. These landmarks functioned as charging stations where the robots could refill their energy supplies by jointly visiting the vicinity of a station.

The robots were equipped with an array of eight equally distanced active infrared sensors surrounding the entire robot. These sensors measured distances between the robot and the nearest object that is in a sensor's visual range. Looking towards the front each robot had a linear colour camera with a resolution of 20×1 pixels, which were distributed horizontally over an angle of 120^o. The colours of the camera were represented by a 2 dimensional vector encoding a normalised representation of a colour's RGB value.[1] The robots' movements were controlled by two independently driven motors. All sensors and actuators were subject to noise, although the colour measured by the cameras were precise.

The control architecture of the robots was based on the schema-based architecture for multi-agent systems as introduced in [2]. It consisted of a finite state automaton (Figure 1) in which each state indicated which schemas were activated (see Table 1). The schemas implemented basic sensorimotor couplings for producing simple behaviours, such as obstacle avoidance and phototaxis (or tracking). In addition, they also implemented higher cognitive functions that were used to produce and interpret linguistic utterances. The combination of these reactive and

[1] By normalising each component of an RGB value to the sum of these three components (Red, Green and Blue), two of these have sufficient information to encode any colour except black. The colour black is treated specially by using negative values.

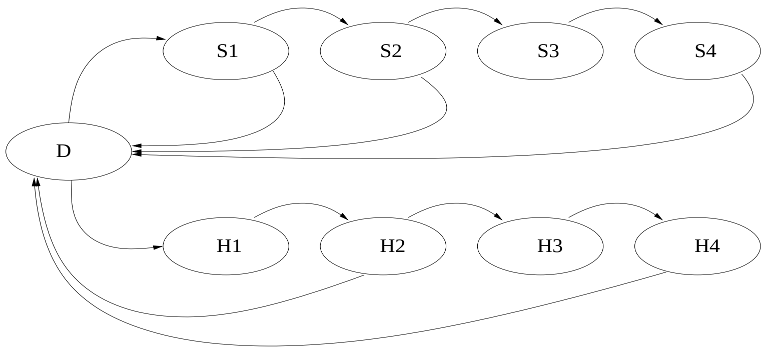

Figure 1: The finite state automaton that guides the sequential behaviour performed by the robots. There is one branch of the automaton regulating the script for the speaker (S) of a language game and another branch regulates the script for the hearer (H).

cognitive schemas justifies to call this a *behaviour-based cognitive architecture*. In the simulation, the finite state automaton (FSA) implemented a *script* for playing a guessing game (see Figure 1).

In a guessing game two robots participated: one speaker and one hearer. The aim of a guessing game was for the speaker to name a meaning (the representation of an action) that the hearer tried to guess based on the uttered name. After performing the action, the robots received feedback concerning the appropriateness of the action. When the proper action was used, the robots reinforced the association between the name and the meaning. And when the action was improper, the association was inhibited. These adaptations are similar to the ones employed by reinforcement learning. When the speaker or the hearer had no name for the action in their lexicons, they would add a new element to their lexicons either by inventing a new name (in case of the speaker) or by adopting the uttered name (in case of the hearer).

At default, both robots explored the environment in an arbitrary fashion using the schemas `Forward` (movement), `OA` (implementing obstacle avoidance) and `Explore`. At particular moments in time a robot could decide to become a speaker and entered state S1 of the FSA.

This state implements the schema `Produce` in which the speaker produces an utterance to indicate the target towards it would go in order to refill its energy supply. This target was selected by determining the most dominant colour in its visual field.[2] The feature vector representation of the target (the 2 dimensional colour vector) was categorised with a prototype that is sufficiently close as modeled in the *identification game* [12]. The identification game is a variant of *instance-based learning* in which the feature vector is categorised using the 1-nearest neighbourhood algorithm with the constriction that the distance between

[2] A colour is most dominant if it occupies the largest subspace in the visual field and if it is not white or black – which are the colours of the walls and robots.

State	D	S1	S2	S3	S4
Schema	Forward OA Explore	Produce	Forward OA Track	Wait	Adapt
Final cond.	S1: Time H1: Expr.	S2: Expr. D: Failure	S3: Target D: Time	S4: Feedback S4: Time	D: Finish
State		H1	H2	H3	H4
Schema		Interpret	Forward OA Track	Wait	Adapt
Final cond.		H2: Goal	H3: Target D: Time	H4: Feedback H4: Time	D: Finish

Table 1: This table shows the implemented FSA (see Figure 1). Each state activates a number of schemas and has one or two final conditions. The final conditions indicate when the agent transits into another state. The following abbreviations are used: D = default state, Sx = speaker's state x, Hx = hearer's state x, OA = obstacle avoidance and Expr. = receiving or uttering an expression.

feature vector and prototype (the category) should not exceed a given threshold. When no such category exists, a new prototype is added to the ontology with the feature vector as an instance. When a category was found, the speaker tried to produce an utterance by selecting that word-meaning association from its lexicon of which the category matched the meaning and the association score was highest.[3] When no such association was found, the speaker invented a new word[4], associated it with the meaning and stored it in its lexicon. The selected word was uttered through the simulated radio communication and the speaker entered the next state S2. The meaning of the utterance was represented by the category of the colour coupled with the activation of the schema `Track` in state S2. When no utterance could be produced, the speaker transferred into the default state D.

At receiving an uttered word, the other robot transferred from the default state into the first hearer state H1. In this state it tried to `Interpret` the uttered word by searching its lexicon for matching elements of which the one with the highest association score was selected. The associated meaning encoded the colour of the target that the hearer guessed the speaker would go to. As is common in reinforcement learning (the technique on which the guessing game is based), a balance has to be sought between exploration and exploitation. This was done using the ϵ-greedy technique, see e.g. [11]. When the hearer did not know the word or when it decided to explore an alternative, it selected the most dominant colour in its visual field as a target. When the hearer thus selected a target, it transferred into state H2 of the FSA.

The states S2 and H2 were similar in that they activated the same schemas

[3]The lexicon consists of elements containing a word, a meaning and an association score. The association scores (or Q-values) indicate the effectiveness of an element based on past interactions.

[4]A word is constructed by combining a number of arbitrary letters taken from the alphabet.

and finished under the same final conditions. The three active schemas (`Forward`, `OA` and `Track`) together implemented the behaviour of phototaxis towards a target (specified by the selected target colour) while avoiding obstacles. When a robot was close to its target it entered state S3 or H3. When a robot failed to find the target within a pre-specified time (`Time`), it returned to the default state D. In that case the guessing game failed. but, in contrast to previous implementations, the lexicon was not adapted because it was unsure whether the communication would have been successful or not.

In states S3 and H3, the robot that arrived first at the target waited until the other robot reached a target as well. If this was the same target, the energy supplies of both robots were refilled, thus providing positive feedback on the guessing game. If it was a different one, a negative reward was received. In either case the robots entered states S4 and H4. If one robot did not reach a target, the waiting was 'timed out' and the waiting robot returned to the default state.

In states S4 and H4, the robots adapted their lexicons according to the feedback they received. If the robots received positive feedback, they reinforced the association score of the used word-meaning association. In addition, both robots laterally inhibited scores of competing associations.[5] If the hearer did not select a target based on the utterance, i.e. the hearer either did not know the word or it decided to explore, then the hearer adopted the word in association with the meaning of the target. (Note that the hearer only adopted a new association when it received a positive response.) When the robots received a negative response, they lowered the association score of the used lexical element (if any was used). When they finished the adaptation, the guessing game finished and the robots returned to the default state.

Summarising, the above framework provided a simulation in which robots tried to coordinate their activities using the guessing game as a communicative guideline. The robots used this guideline to select a target charging station. If either robot failed to use the communication to select a target, this robot selected an arbitrary target. According to the effect of the communication, the robots were able to adapt their lexicons in the same way as in previous implementations of the guessing game, e.g. [10, 13]. Because the thus emerging lexicon was used to improve their ability to refill their energy supplies, this lexicon was used in a functional manner. The representation of a word's meaning was based on the representation of the target (i.e. its colour) and the activation of the proper behaviour (i.e. the schema `Track`), rather than on the representations of the behaviours' sensorimotor data.

3 Results

In this section the results of two simulations are presented. In one condition, no communication was used, but when a guessing game started, both robots selected the target colour that was most dominat in their visual fields. This condition was tested to see what level of success the robots could reach without using language.

[5] An association is competing when either its word matches the used word, but not its meaning, or when its meaning matches the used meaning, but not its word.

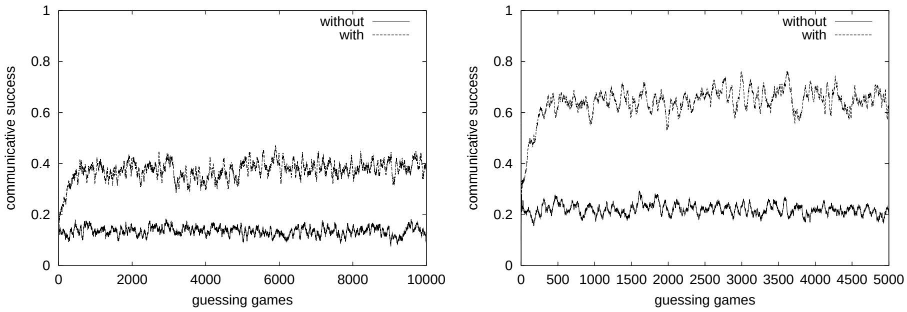

Figure 2: The results of the experiments. The figures show the communicative success as a function of the number of guessing games played (left) and as a function of *finished* guessing games (right). Each figure shows the results of experiments 'with' and 'without' communication.

The other condition was exactly as described in the previous section. Both conditions were simulated 10 times under different random seeds and the measured results were averaged over these 10 runs.

Figure 2 (left) shows the communicative success of both simulations as a function of the number of guessing games that the robots played.[6] It is clear from this graph that the simulations with communication outperformed the simulations in which no communication is used. Without communication the success fluctuated around 15%. The communicative success when communication was used stabilised around an average of 40 % from approximately 1000 games. When inspecting the communicative success as a function of the number of games in which the guessing game finished[7], the level of communicative success increased to a value around 70 % with communication (Figure 2 (right)). This is again much higher than the success reached without using communication, which was approximately 20 %.

4 Discussion

The relative simplicity of the simulation's setup may induce the suggestion that the research is trivial, but it appears not so. Given that simulations with the guessing game under, apparently, similar ideal circumstances yielded perfect results [7], also similar high results were expected in the current simulation. However, the results did not exceed a level of 70 % in communication accuracy, which is not higher than results obtained in experiments using real robots [3, 13]. A possible candidate for the imperfect communication is that the robots decide they are near the target while they are near another landmark. The robots decide that they are near the target charging station when the front infrared sensor senses it is near an

[6]The communicative success measures the number of games in which both robots refilled their energy supplies, averaged over the past 50 games.

[7]A guessing game finishes when both robots reached a target.

object and when the centre of the visual field detects the searched target colour. However, because the resolution of the infrared sensor is much smaller than the camera's resolution, a robot may decide it is near the target when it is actually near something else while 'seeing' the target. If this actually happens has to be checked in future work.

The main purpose of this paper is to illustrate how the meaning of single word utterances can be represented rather easily on top of already existing reactive behaviours. Many roboticists, including myself, have anchored the words (or symbols) to the sensorimotor data representing actions [5, 8, 12]. But the simulations show that rather complex meanings can be represented by coupling a target colour to the control module (or schema) for moving somewhere. The representation of the utterances' meaning recruit already existing schemas, thus constructing the language on top of these evolutionary older structures.

To illustrate the complexity of the meanings let us analyse a typical guessing game. Suppose that the speaker selects the red target as the station it will home in on. The utterance it produces, for instance "wabako", then means something like ¡I will go to the red target¿. For the hearer, when it interprets the utterance and takes the appropriate response, the utterance may mean something like ¡I think the speaker goes to the red target, and so will I¿. This is perhaps a bit overdone as the hearer has no notion what the speaker is doing, nor does it have a notion what it is doing. But implicitly such a meaning may arise and can certainly be addressed from an observer's point of view. And when the communication is used in relation to a 'life-task' that the robots must perform in order to 'stay viable over extended periods of time', the utterance might really become meaningful to a robot [15].

In the current simulation the words are only associated with prototypes representing a colour, coupled with the activation of the schema `Target` that implements the reactive sensorimotor control for navigating towards a target colour. An interesting extension would be when the robots have the ability to select between various schemas that implement different behaviours. This would allow robots to construct a much richer repertoire of meanings describing actions.

With such an extension it is also possible to investigate whether the robots can also learn the proper behaviour from another robot. An interesting experiment would be one in which the robots have to select whether they couple a target colour with a behaviour that attracts them towards or that repels them from a target. Such an experiment can even be used to enable robots to develop a proper set of coordinated behaviours in order to "survive" over extended periods of time as shown without communication in, e.g., [9]. This way the robots construct symbolic communication with respect to their "life-task", which might well be necessary to form symbols that are really meaningful to the robots themselves [15].

References

[1] R. C. Arkin. Motor schema-based mobile robot navigation. *International Journal of Robotics Research*, 8(4):92–112, 1989.

[2] T. Balch and R.C. Arkin. Communication in reactive multiagent robotic systems. *Autonomous Robots*, 1(1):27–52, 1995.

[3] A. Billard and K. Dautenhahn. Experiments in social robotics: grounding and use of communication in autonomous agents. *Adaptive Behavior*, 7(3-4):415–438, 1999.

[4] R. A. Brooks. Elephants don't play chess. *Robotics and Autonomous Systems*, 6:3–15, 1990.

[5] S. Coradeschi and A. Saffiotti. Anchoring symbols to sensor data: preliminary report. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence (AAAI-2000)*, pages 129–135, Austin, 2000.

[6] S. Harnad. The symbol grounding problem. *Physica D*, 42:335–346, 1990.

[7] F. Kaplan. *La naissance d' une langue chez les robots.* Hermes Science, 2001.

[8] M. Rosenstein and P. R. Cohen. Concepts from time series. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*, Menlo Park Ca., 1998. AAAI Press.

[9] L. Steels. Discovering the competitors. *Adaptive Behavior*, 4(2):173–199, 1996.

[10] L. Steels and F. Kaplan. Situated grounded word semantics. In *Proceedings of IJCAI 99.* Morgan Kaufmann, 1999.

[11] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction.* The MIT Press (A Bratford Book), Cambridge, MA, 1998.

[12] P. Vogt. Grounding language about actions: Mobile robots playing follow me games. In Meyer, Bertholz, Floreano, Roitblat, and Wilson, editors, *SAB2000 Proceedings Supplement Book*, Honolulu, 2000. International Society for Adaptive Behavior.

[13] P. Vogt. The physical symbol grounding problem. *Cognitive Systems Research Journal, special issue on Situated and Embodied Cognition*, 2002. in press.

[14] H. Yanco and L. Stein. An adaptive communication protocol for cooperating mobile robots. In J-A. Meyer, H. L. Roitblat, and S. Wilson, editors, *From Animals to Animats 2. Proceedings of the Second International Conference on Simulation of Adaptive Behavior*, pages 478–485, Cambridge Ma., 1993. The MIT Press.

[15] T. Ziemke and N. E. Sharkey. A stroll through the worlds of robots and animals: Applying jakob von uexküll's theory of meaning to adaptive robots and artificial life. *Semiotica*, 134(1-4):701–746, 2001.