Chapter 4

Integrating Talk with Action

Introduction

In the previous chapter we considered existing work to develop signalling behaviour within the bottom up approach to artificial Intelligence. It was shown that this field started with quite general experiments, establishing the scope of the discipline and then broke into a few separate branches to investigate specific features of signalling and linguistics, though the boundaries between them were in many cases quite indistinct. By having reviewed the field in this manner it should now be possible to establish where further investigation is required if we are to work towards understanding how agents might develop complex linguistic capabilities and any of the additional benefits that might be associated with this.

We shall now consider an argument for why there needs to be a greater understanding of how evolving agents might achieve the integration of the production and reception of signals with other developing behaviours. Once this is established, potential difficulties that have restricted work within this discipline up until now are exposed, from which it will be possible to propose some empirical questions and an experimental framework in which ways to overcome these hurdles can be investigated. This proposed direction of scientific enquiry is made up of the original ideas that were to be explored by this research project and objectives of the actual experiments that were conducted. The two differ on account of difficulties arising in the earlier experiments that narrowed the field of enquiry. An overview of these experiments will be discussed in preparation for the full details laid out in the following chapters.

4.1 Integration of Signals to Sensors and Behaviour

Whilst some of the work that we reviewed was quite abstract, particularly many of the investigations considering the adaptability of grammars, there were also a number of examples where signalling issues were considered with respect to real robotic problems. Leaving aside issues relating to signal motivation for the moment, the majority of these projects were concerned with integration of developing signalling systems with the parts of the agents that connected them with their environments, notably their sensors for input and behaviour actuators. Typically, agents would have to adapt to adopt the role of signal receiver or transmitter. In a number of experiments these were mutually exclusive whereas in others an agent could take either role. An agent in the role of signal transmitter would determine the state of its immediate environment through its sensors, and accordingly produce either a single signal or some more complex combination of signals. The agent designated the role of signal receiver in this interaction would detect the signal with its signal sensors and through its control system determine a particular behaviour to be activated via its outputs. The agents were required to coordinate the activation of certain behaviours predicated on the detection by others of specific environmental features. This required Saussurean signalling systems which meant one to one mappings between sensor states and behaviour activated through a coordinated signalling system.

A major limitation on these experiments was the restricted range of states that could effect, and be the subject of, an act of communication and also the limited range of behaviours that could result from a signalling interaction. At one extreme we might consider the work of Oliphant[1996] where agents might each have a single sensor, that can take one of two states, the state of which they are required to communicate. The agents in the role of receivers in this communication interaction have only a single simple behaviour that is also restricted to two states. Such simplicity of signalling system is to be expected in many of the investigations to understand motivations for cooperation required for signalling but almost as simple models of environmental interaction occurred in the more complex experiments that we reviewed of Steels and Cangelosi. For example, in Cangelosi's experiment to generate symbolic signalling,[1999] although the behaviours and sensors seemed complex, most of the complexity appears from the interpretations of the behaviours and sensors such as 'mushroom feature' detector and 'eat mushroom behaviour' actuator that are imposed upon the model to provide it with some contextualisation. In this experiment the agents have only a few behaviours and these are independent and not made up of component parts. This was a way to make the experiments more tractable as we shall see later, so that high-level features relating to the development of symbolic signals could be investigated. Similarly, Steels' experiments with the talking heads robots while dealing with a very complex problem domain are designed in such a way that the adaptable communication systems within the architecture was also presented with a very simplified part of the problem.

In these experiments, what is deemed to be the development of signalling, is the coordination of the signal transmission and reception such that there can be a one to one mapping between environmental sensors[1] of the transmitter and behaviours of the receivers. Within these experiments there are a fixed number of environmental cues for signalling and the restricted behaviours associated with each output state are the only possible outcomes of an interaction.

Whilst this approach allow one to get to grips with the issues of how symbol coordination can occur, it requires a comprehensive prior understanding of the problem space to be able to furnish the agents with appropriate sensors and behaviours. This means that analysis has to be performed on the task to be completed by the agents (with the aid of signalling) so that the appropriate cues have corresponding sensors feeding into the signal transmission system and the range of responses required by the agents receiving signals must be determined and provided for the agents. There are two problems with this approach.

In the first case, using the adaptive bottom up approach to artificial intelligence to create signalling strategies has been partly championed because we want to increase the burden of choice of behaviours, feature categorisations and appropriate strategies onto the adaptive system rather than leaving this to a human designer. Clearly at some level the designer must choose a set of sensors and base behaviours, but the lower level these are, so we get an increase in their generality and range of uses they can be put to with an appropriately adapted control system. It is also necessary that these lower level behaviours would have to be able to combine in such a way to create behaviours suitable for the task and so too would the control architecture have to be capable of combining information from feature sensors to determine higher order features within the environment that might be a suitable cue for signalling. Whilst an agent furnished with such low level appropriately selected base sensors and primitive behaviours would have a harder task adapting to higher level problems, it would have a greatly increased range of problems that it could adapt to.

---

[1] Here, environmental could well refer to inner states of the agent where for example, the agent might be communicating about the state of its *hunger* or *strength.*

The second problem with restricting the repetoire of environmental interactions an agent can perform whilst engaging in a communicative task is that this will act to restrict the expressive nature of the communication capabilities of the agents. Let us consider a crude example to make this point clear. We shall consider two separate types of agents, the first have modular behaviours corresponding to a particular signalling task whilst the second have slightly lower level behaviours that need to be combined to perform the necessary behaviour to take advantage of a communicative act.

1. Agents have a behaviour they can activate to return to *battery charging base*. There is some environmental feature that allows an agent to determine whether their partner's batteries are running down. Agents can employ a single signal so that when one detects that its partner's batteries are low, the partner can return to the charger to refresh. The signalling system is very straight forward requiring a single connection for the signal issuing agent that connects between the low partner energy sensor and a free signal channel, and for the signal receiving agent a single connection is required between the corresponding signal receptor and the *return to base* actuator.

2. Agents have a series of component behaviours, Move in a straight line, turn towards left and turn towards right These agents also have a similar sensor as the agents in the previous example to detect their partners energy level. Furthermore, they have sensors to determine the position of their partner relative to the *battery charging base*[2]. In order for these agents to achieve the same feat of assisting each other to remain charged, the level of communication will have to be more sophisticated with one agent not simply telling the other to go and charge, but also having to tell it where and how to move in order to do this.

Whilst the development of an effective signalling system in the second example would be a considerably harder feat to achieve, it should be clear that this second approach offers more flexibility and thus a greater chance of exploring a developing communication system. This is because the combinations of behaviours are less restricted. Therefore, where there arise minor variations in the tasks that the agents are to achieve, slight variations in the signalling can be explored allowing for an incremental understanding of communication development. By allowing the signalling system to integrate with the agents' controls at almost an *atomic* level in

[2] This does not have to be a complex feature, possibly just combining the directions of the two objects relative to itself.

terms of sensor states and actions, we offer the possibility of a communication system that can result in the production of any behaviour available to the agents and convey information about or be influenced by any feature within their environment that they are able to discriminate. Within behavioural robotics we see behaviours expressed down to the levels of simple motor actions and primitive sensors at the level of range finders and pixels representing light levels from visual arrays within cameras. We don't want to stop out integration at this *atomic* level with signalling restricted to interfacing only with raw sensors and actuators. The signalling system should have access to all areas of an agent's control system, the potentially higher level information coming from within the control systems[3] as well as individual variables that an agent can modify to activate some complex series of behaviours.

By increasing the level of complexity at which a signalling system interacts with the rest of an agent's control system we assist in prescribing how the task should be performed but this is done at the expense of increasing the limitations on what uses the signalling system can be put to by the agents. One of the qualities we ascribe to human language that gives it such immense power and value is its seemingly open nature, in that it can be used to refer to anything we are able to sense and can be used to activate virtually any behaviour we are capable of. This must surely be a quality that we are to strive for with the signalling systems of our artificial agents, signals capturing the most subtle variations in a particular environment or generating an infinite variety of responses from a listener. Now our difficulty will be to allow this to happen within agent populations. Later in the chapter we shall see why the problem with this approach is considerably harder than the impression of a vast task one initially gets and this accounts for why the experiments we have seen so far start off at such high levels of information and control even given that nearly all of the researchers we have considered are embarking on a bottom up project.

4.2 Dedicated Channels?

At this stage it would be fair to wonder if this issue is exactly the problem that is being addressed by Quinn in his experiments [2001] that we reviewed at the end of the last chapter. Quinn has demonstrated how signalling can arise without the need for dedicated signalling

[3] For example, one part of a robot control might of adapted to detect mushrooms, relying on some finer grained sensors such as a visual array. A particular state within this part of the control system could be activated when a mushroom is within a certain range.

channels within agent populations with the sort of low level sensors and actuators that we have just discussed. Furthermore, he believes that the signalling behaviour itself must derive from these same low level sensors and actuators. He has argued that the absence of such dedicated channels is a necessary position to take if we are to use an evolutionary modelling approach to help better understand origins of communication behaviour as it seems that much animal signalling can be explained as an adoption of a pre-existing behaviour or trait:

> "Open any textbook which gives an evolutionary account of animal behaviour and you will find hypothetical reconstructions of the processes surrounding the origins and early evolution of communicative behaviour which revolve around non-communicative behaviours acquiring selective value in a communicative context" p.10 Quinn 2001

Quinn provides the examples of Hinde [1966], Krebs & Davies[1981] and Maynard Smith[1997] to support this. Whilst Quinn believes that dedicated pre-formed signalling channels are a hinderance to understanding origins of communication, he accepts that such channels have their use for models that already have communication in place but wish to explore this further

> "Dedicated channels are a reasonable feature of a model which assumes that individuals are already able to communicate."p.1 Quinn 2001

We shall now consider an argument for why the exploitation of dedicated channels within agent based modelling for the enhancement of our understanding of evolutionary biology and 'hard' artificial intelligence is not without value. We shall also consider why this same argument suggests that the, all be it un-natural, provision of dedicated channels to artificial agents could possibly assist in the development of linguistic capabilities that might take a considerable longer time to develop within a system where signals are derived from pre-existing behaviour incorporating the transmission of information without a communicative function.

We can see preformed dedicated channels as any conduit that is capable of conveying information and serves no other function within the particular model. Quite how does this differ from channels that evolve around the non-communicative behaviours referred to by Quinn? The latter exist because it is possible for certain information to be transmitted and by behaviours within the agents' behavioural repetoires that are intended to serve primarily for non-communicative tasks. Also, the agents that would be potential receivers of this information would be required to have sensors capable of detecting and discriminating the information from other environmental noise again for non communicative tasks. Referring back to his experiment, Quinn exploited proximity detectors that the agents use for navigation and simply the motion of the agents as the information. With this approach we see behaviour originally not being used for communication but to some extent achieving this, being modified to assist in the information transfer function. This modification produces a resultant behaviour to better accommodate the original function as well as the communicative task.

There is a limitation to the signalling behaviour that can arise with this method. By their nature the signals will be of an iconic nature. It is not possible for symbolic signals to arise as any signal derives itself directly from the action that it represents. With these iconic signals, it is necessary for the behaviour to be detected by the receiver. In many problem domains that could benefit from signalling it could be argued that there might be some subtle information processing happening within the agent from its interpretation of environmental features and it then behaves in a manner that is not possible for the potential receiver of the signal to interpret. The behaviours could be quite different but simply not within the capabilities of the receiver to interpret them. This is almost like the reverse of having the dedicated signalling channel because in these cases a signal is being issued but one that is not available to the receiver. It is conceivable that this will not only arise in a very limited set of circumstances particularly if one is to believe that a significant number of communicative acts serve as information transfer where one agent might have information not available to another. Under these conditions it is likely that some physical separation exists where limited information is being transferred except specifically as part of an act of communication.

By relying on dedicated channels for the formation of a signalling system we can avoid some of these problems. With this approach, there is an arbitrary connection between behaviour of the potential signal issuer and the signal production so that a signal is not connected to the behaviour of the potential signal transmitter but rather to any part of the agent that reflects the state of information to be transferred. Similarly, for the signal receiver there is an arbitrary connection between the signal channel and its behaviour. Provided there is sufficient information on the channel then it should be possible for the receiver to discriminate this and modulate its behaviour accordingly. This arbitrary connection at both the transmitter and receivers side of the signalling system opens the way for symbolic signalling.

But, does having a dedicated channel in the first place mean that we have failed to explain how signalling arises? By taking advantage of dedicated signal channels, models do not have to be claiming that the channel exists exclusively for the purpose of communication, or the communication of the particular experiment for which they are being used. In fact the assumption that they are making is that the *dedicated channel* exists for some other function, possibly not communicative but also for some other communication function, but importantly it has the function of being able to transfer information between two or more agents and is thus available to be used for such a function. However, if the function for which the information conduit originally exists is irrelevant to the experiment then there is simply no need to model this aspect of its functionality within the particular experiment. For example, the functionality of communication could be of considerably greater importance to the original function thereby completely modifying its function or the original function might exist with no interference to or from the communicative use of the channel and be completely unrelated to any of the communicative behaviour. Clearly any evolutionary model is an abstraction and it should not be necessary and more importantly is not possible, to include all characteristics of the types of animals that we might be modelling the evolution of.

There will be occasions where we shall be interested in both approaches. A hybrid system is possible where some signals derive from non-signalling behaviours whilst employing other signals that exploit the signal channel in an arbitrary fashion. The experiment in the next chapter uses a channel that has derived functionality with respect of the signal receivers but it is a dedicated signal channel for the transmitters. The reason for this hybrid approach will be discussed in the next section. Having argued for the advantage of using dedicated signalling channels we shall now take a look at a big disadvantage that they offer.

4.3 Four Part Problem

I suggest that it is possible for the acquisition of symbolic signalling behaviour integrated with other low level behaviours but there are substantial problems developing this. With this thesis we wish to investigate how it might be possible to overcome these issues in such a way that is possible for our artificial agents but might not have been the case in natural systems. We have just considered why dedicated channels might offer an advantage over signalling systems that arise through iconic signalling evolving as adaptations of existing behaviours but just the inclusion of preformed dedicated channels integrated at a low level is clearly not enough or else we can assume that this approach would have already been explored and as we have seen in our review, experiments have avoided this in a number of ways. Four such ways are summarised.

<u>High level abstraction</u>

Where we have seen dedicated signal channels used in models attempting to develop signalling systems within populations of adaptive agents, we have nearly always encountered quite severe restrictions imposed on the potential complexity of the signalling system to be developed. Earlier in this chapter we saw how some experiments within this problem domain relied on a very simplified agent architectures that only provided potential signalling systems with access to high level sensors and behaviours, and as a result were missing a lot of the advantages that we would observe if these higher level capabilities were formed from low level behaviours and sensors integrated with the signalling behaviour. Whilst there was an advantage of this higher level approach in terms of exposition of ideas by keeping the models as simple as need be for the purposes they were trying to demonstrate the disadvantages will manifest themselves as these systems are expanded to incorporate more complex signalling behaviour.

<u>Restricted architectures</u>

In a number of the experiments that we reviewed, for example Oliphant[] and Nobel[], the focus of the experiments was to understand the motivations for communication. As a result, the agents were kept as simple as possible. Even if one were to argue that these experiments do not fit within the previous category of high level abstactions, for similar reasons to those experiments the number of possible control configurations is so small that we are only ever going to see very restricted signalling behaviours emerging in these models.

<u>Sculpted fitness function</u>

Saunders and Pollack[1996] were able to show that it was possible for signalling to arise in a relatively complex environment where a learning mechanism exists which isolates the various component parts of the signalling problem and encourages the isolated development of each of these parts. The signalling system arising once all component parts are in place. This way of developing signalling requires the experimenter to specify the general outline of the communication behaviour to be used with the adaptive approach providing the implementation. We considered in 4.1 problems with this approach for our project.

<u>Partial provision of a solution</u>

In the work of Billard and Dautenhahn [1997] we have seen that the problem space is reduced by the provision of a partial solution to the problem. Here, though relatively low level behaviours were available for two robots exploring a landscape with one robot with the responsibility of following the other, the leader robot was provided with a signalling system which provided a *description* of its location on account of a series of signals equating with geographical features of the environment. Here the signalling task was simply for the follower to develop a receiving strategy around this preconceived signal transmission behaviour.

At the outset we claimed that the development of a signalling system was a difficult problem and the noticeable absence of symbolic signalling systems besides our own language appearing within the natural world [Deacon 1997] alerts us to this problem. Within the situated adaptive behaviour approach to artificial intelligence we are able to simplify the task of developing signalling by isolating the issues of motivation for the signalling by providing tasks in which such signalling behaviour is rewarded with no conflict[4]. So why does all work even where agent motivation is catered for need to take the short cuts referred to above and reduce the overall generality and power of their resulting signalling system? Why don't we already have communication systems that employ low level behaviours that are already capable of exploiting complex syntactic features especially given that these were stated objectives for following up on the original works within this domain nearly a decade ago [Werner & Dyer 1992] [McLennan 1992]. Clearly there is some problem that all these approaches are trying to avoid. This problem we shall refer to as the four part coordination problem which we shall now look at.

The difficulty stems from the number of different and possibly independent features that have to arise and be coordinated in order for any benefits to accrue from signalling behaviour.

[4] Examples of conflicts that might appear in nature would be the costs associated with energy required for signal production or the increased risk of predation on account of the signals being used attracting more predators.

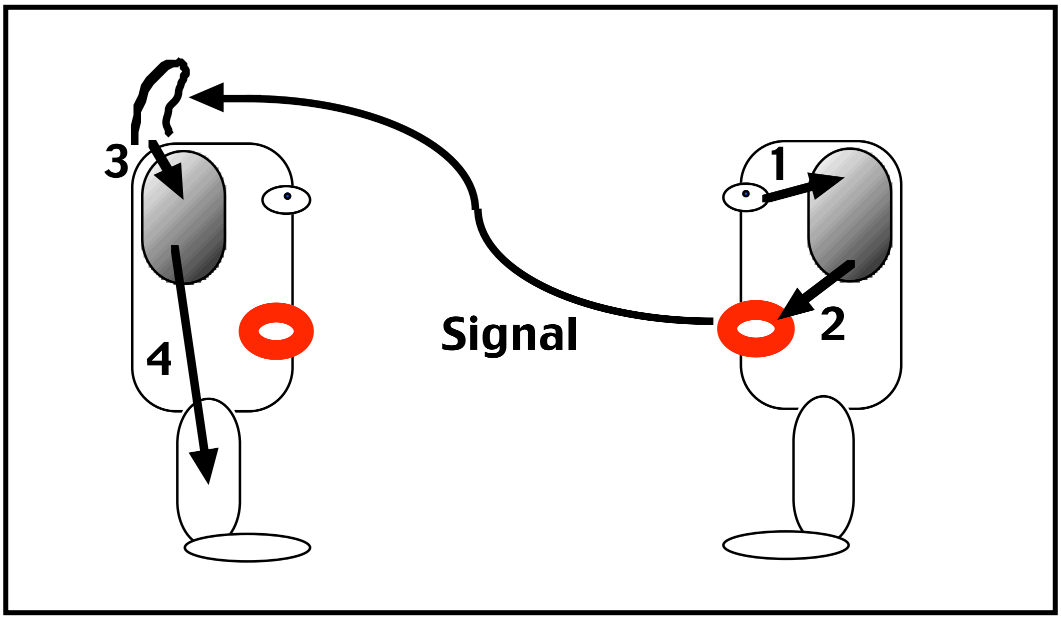

Figure 4.1 *Four independent stages required for successful communication behaviour.*

Consider we have two agents involved in an interaction from which both will benefit if one (*T*) conveys some information to the other (*R*) and this second agent acts appropriately to the value of the information. The state of this information only agent *T* can potentially retrieve from either the environment or its inner state. In order for these two to benefit

1 *T* must detect and discriminate this situation that warrants a signal and then
2 activate its signal transmitter device to create a signal unique to the situation.
3 *R* must detect the signal and discriminate it from others and finally
4 *R* needs to modify its behaviour accordingly to take advantage of this new information.

These stages are referred to in figure 4.1.

There is a multiplicative increase in search space of possible control architectures when one wishes to combine these separate systems and this is what makes signalling such a hard task. As the size of these separate spaces increases, the size of the overall problem increases in a combinatorial manner. Where there are a small number of possible world states that agents are able to sense and a very restricted number of possible behavioural repertoires and signals they can generate relating to an interaction, then the search space of suitable control systems is not too large. As we have just seen, most existing work within this domain has applied this approach with a view to understanding higher level characteristics of signalling systems. However, it seems intuitive that more general communicative abilities with the possibility of increased complexity to cover non-finite problem domains (including most real world problems) must be integrated at a very low level of control capable of generating an infinite range of behaviours. But it is just this approach where the four separate parts mentioned above each present a very big search space in their own right and unfeasibly large combined search space. Allowing the four separate stages to develop in their own way and coincidentally match at some point would simply take too long, therefore we need to explore ways of speeding up the process. The rest of this thesis concerns itself with an exploration of ways in which we can overcome this four part coordination problem. We shall now consider how experimental work will be used for this exploration.

4.4 Empirical Approach / Solutions

In order to understand how we might assist agents in developing symbolic signalling behaviour that is integrated with low level control features available to the agents, we obviously require agents with such design. Furthermore, we shall be looking at how the agents acquire the signalling behaviour as part of the solution to some higher level task. By selecting a task commonly encountered within the situated robotics research literature and also one that is of significant importance to many robotics tasks, any understanding of how the signalling develops will have a significant use even if the lessons learned do not have a direct use within other problem domains. However, we have striven to keep the principles learned about signalling as general as possible and thus applicable to many other problem domains.

We start by defining an environment in which these experiments can exist. The point of this is to work within existing limits of what has been achieved with bottom up artificial intelligence but also to provide an environment in which it is possible to provide an incremental augmentation of complexity to allow for a corresponding augmentation of the complexity of the signalling behaviour. We limit the problem domain to a navigation task, performed by simulated mobile robots, that requires complexity of coordination of various sensors and motor actions. To provide an arena that would promote the development of communication behaviour, several robots are introduced to the maze arena in such a way to allow them to share the task of exploration through the use of communication protocols that they might develop. Initially some consideration went to providing a complex social environment relying on the principles of machiavellian intelligence [Byrne & Whitten] which assumes that the most complex environmental features come from interaction with other agents, but this was rejected as an arena for investigating communication as it introduced too many additional variables and furthermore, signals produced in such an environment might be very subtle and difficult to detect.

Three particular approaches to overcoming the four part coordination problem are considered separately and each of these takes up one of the next three chapters. These evolved rather than being planned in advance. The reason for this approach was simply that the original experiments proved harder to generate signalling than was initially envisaged. Difficulties with each of the sets of the experiments led to the subsequent experiments. The final experiments opened up yet more questions left for further investigation. The three approaches to acquiring low level symbolic signalling are as follows:

1 imitation;
2 babbling; and
3 thinking aloud.

In the first of these approaches we explore how agents can use *imitation* to exploit pre-existing signals within their environment. We investigate how signalling can arise by agents adapting to exploit some functionality of one half of the potential communication system, the reception behaviour, that has previously developed to make use of signals generated from their environment. What we shall see is a transmitting system that develops to imitate existing environmental signals and serves no other function than to integrate with the previously formed reception behaviour. Once this restricted signalling is shown to develop we then consider how the signalling behaviour can develop further from this state both in terms of the stability and sustainability of the system even if the original source of signals from the environment is removed. The enhancement of this signalling system for the acquisition of new signals is also considered.

The original goal of the research program was to spend more time with the model we have just discussed, looking to see how as new signals were derived, they would reflect with a syntactic structure, the structure of the environment in which they are used, possibly with formation of simple numbers and spatial relationships, useful for giving increasingly more complex instructions. Unfortunately, the difficulty at this stage in creating multi signal systems turned my attention to understanding what various effects of the agent architecture would allow for better simple signal acquisition. This led to the consideration of two further approaches to creating signalling behaviour within adapting agents. The first of these we shall refer to as the *babbling* approach and the second is achieved by *thinking aloud*. As we shall see, each of the experiments tries to overcome the difficulty of combining several complex search spaces from different sides. The first that we have mentioned, considers some of the issues when the signalling has to from around some preformed reception behaviour. The further two sets of experiments consider how the chance of signalling can be enhanced without taking advantage of this relevant non-communicative behaviour besides the provision of dedicated signalling channels.

The babbling approach explores what are the factors that effect the rate of signal acquisition if we provide a control architecture for the agents that forces upon them the simultaneous production of many signals to attempt to increase the rate of exploration of the search space of possible signalling strategies. This approach limits the size of the search space by restricting it to architectures that automatically generate signals. Most of these are of no relevance to the problem domain, but by keeping this highly dynamic, we wish to see whether signalling behaviour can occur and if it does whether the signalling acts to reduce the high variation in signalling strategies that was used to generate the signalling but could equally destroy any such strategy as well.

A refined version of the babbling approach is then considered which is suited for a subset of possible signalling tasks (albeit a very large subset). This *talking aloud* version relies on the agents optimising the chance of producing relevant signals with respect of their problem domain by employing a control architecture that requires the agents to produce externalised signals for their own control. For many tasks of communication the information that needs to be communicated also plays some part in the non-communicative behaviour of the signal issuer. In these situations, this approach tackles the complexity of search space by allowing the transmission side of the behaviour to develop first. Whilst this might look like we are returning to the idea that signalling must derive from non-dedicated channels, what we are doing is changing the emphasis so that such signal channels are necessary for other behaviours and this explicitly forces the sharing of information that might not ever be possible to discriminate using normal sensors and behaviours.