Chapter 8

Conclusions

This thesis set out to understand how signalling behaviour could be developed within the evolutionary robotics framework such that it would be a suitable foundation from which future efforts to generate linguistic behaviour could be based. Given the difficulties that arise within the fields of linguistics and evolutionary biology, on account of the paucity of historical data, to generate very strong evidence for any particular account for the origins of language, this is a very exploratory discipline. Though the work of this thesis was intended solely for the enhancement of a bottom up approach to artificial intelligence, it is possible that it will ultimately contribute to either of the afore mentioned disciplines by working to increase the number of examples of signalling and thus help to define the boundaries of signalling behaviour in general.

This chapter summarises the major contributions made by the thesis. These are:

- An argument is made for the necessity of bottom up AI to acquire the ability to manipulate grounded symbols through the acquisition of communication behaviour;
- A categorisation of existing work within the discipline is made, identifying a number of logical sub-disciplines;
- A case is made for the importance of integrating signalling behaviour with base level sensors and actuators in order for the possibility of developing open language.

Then using a modelled population of evolving robots the following original results are obtained:

- Imitation behaviour is shown to act as a source for dynamically stable signalling systems, even where the original imitated source is removed;
- The constraints imposed by transmission and reception strategies of communicating agents ensure that there is a non-zero finite optimum number of channels for agents to employ;
- It is shown that using a control architecture that forces information passing between sensors and actuators to be externalised can significantly increase the rate at which signalling can be acquired for a range of problem domains.

An argument was presented in Chapter 2 for why evolutionary robotics would benefit from the development of linguistic behaviour. Besides the obvious tangible advantages bestowed upon agents capable of exploiting language such as coordination for collective tasks and increased rates of learning brought about from an 'extended mind', an idea championed by Clark & Chalmers[1997], a case was made for the necessity of language development as a means of introducing into bottom-up AI many of the advantages developed with the physical symbol system, possibly a grounded version of Fodor's language of thought [1975]. Continuation of the work undertaken in this thesis might eventually lead to the confirmation of these ideas.

A comprehensive review of existing work relating to signalling and language development from a bottom up, situated perspective was undertaken in chapter 3. After initial exploratory experiments to establish the discipline undertaken by Werner & Dyer, and MacLennan, work fell predominantly into three categories. These addressed both cooperating and defecting conditions for benefiting from signalling behaviour, the way in which higher linguistic capabilities could develop from a signalling substrate, and the integration of signalling behaviour with low level sensors, controls and behaviours. The final category was pursued in chapter 4. Existing work within this category had relied on signalling being derived from sensors and behaviours required for other modalities of behaviour besides signalling, in fact rejecting sensors and actuators provided explicitly for the purpose of the agents using dedicated signalling channels. But an argument was presented to support the use of dedicated channels in the pursuit of strong AI. One of the consequences of attempting to exploit dedicated channels was a great increase in the difficulty of creating both reception and transmission strategies that would be integrated. This was suggested to be a major stumbling block in the acquisition of symbolic signalling so three possible ways to increase the chances of developing integrated strategies were suggested. These were:

1) to extend imitation behaviour,;
2) babbling and
3) forced externalisation of thought.

The subsequent three chapters evaluated each of these in turn through the use of populations of simulated Khepera robots being evolved to perform cooperative navigation tasks that required partner agents to share information about the position of resources in a maze environment.

It was possible to show that imitation of existing environmental signals offered a relatively easy way for signalling to evolve by providing an incentive for one half of the communication system to develop, the reception strategies, in isolation which would in turn offer a greater opportunity for coordinated transmission behaviour to arise. Signalling generated in this manner proved stable even when the source of imitation was removed but development of the signalling through a memetic approach proved very difficult to achieve though an example in which a further signal was derived suggests that this approach could offer some future hope. One of the major drawbacks with deriving signalling from imitation is the need to provide an original source to be imitated. It can be argued that this is akin to teaching a child an existing language but for many tasks such investment in training might not be desired and furthermore, what we as humans might chose to be suitable categories to signal about for a particular task might not be best suited to the control structure of the agent population to undertake the task.

To avoid imitation, a way was needed to increase the chances of signal transmission and reception behaviour arising and being coordinated simultaneously in a population's history. In Chapter 6, behaviour based on babbling was considered. I t was hypothesised that signalling would evolve more rapidly by increasing the number of signal channels available to the agents above the minimum required to achieve the task. This was demonstrated to be true but the rate at which signalling arose for even a very simple task rose significantly over performances with the aid of imitation. Prior to the development of a successful signalling strategy, very long periods of population fitness stability would be observed where no coordinated transmission and reception strategy could exist simultaneously. Once both were present, each would act to maintain the stability of the other, but with only one arising, there was no evolutionary force to maintain it and there was no guarantee that it would be maintained for a long period. Decreasing the rate of variation between generations would help maintain a particular strategy for longer but at the expense of reducing the number of matching strategies tried against it each generation.

Given that for many circumstances under which signalling would be beneficial, the behaviour of not only the potential receiver but also the potential transmitter of a signal is affected by the information that needs to be transmitted it was argued that by making the potential transmitters express their own control signals by having to resort to external signals to pass information between the sensor side and motor side of their control systems, appropriate stable transmission behaviour would arise within many problem domains from just the fitness rewards accrued from their non-signal related behaviour. This was demonstrated in chapter 7. For similar tasks undertaken by the *quiet thinkers* of the previous chapter, the agents in chapter 7 who had to rely on *externalised though* achieved coordinated signalling at a consistently faster rate. A potential problem of exploiting this externalised thought process is that for some problem domains, the forced exposure of information could be detrimental to the individual. It would be interesting to see agents adapt to whisper such information. This would rely on a more complex environment than was employed in this thesis, with signals decaying over distance.

Each of the three sets of experiments dealt with a very limited problem domain where only a single signal or at most 2 signals were generated. Unfortunately this was due to the large computational burden of the experiments with many taking several weeks each to complete. Furthermore, many initial trial experiments had to be conducted in an attempt to select optimal performance values for the hundreds of parameters used to create the models. In order to extend the experiments to deal with multiple signals and the possibility of syntactic relationships between them a way of increasing the computational performance would be required. Whilst it is possible that minimal simulation techniques [Jakobi 1997] could go some way to assisting with this problem, it is likely that significant increases in computational resources through massive parallelism and increased processor speed will make it easier to extend this research.

One final application for the discovery of the effectiveness of the thinking aloud strategy could be employed in inter-agent component communication. If a robot is made of many adaptive components that are each developed to some extent in isolation to ensure the acquisition of specific behaviours, then by allowing the components to use a universally (within the robot) shared data bus (akin to the dedicated signal channel) to pass information between the sensors and motor side of each component then for a range of problem domains this could increase the likelihood of useful information being shared by components.