Robots As Socially Intelligent Agents

Ruth Aylett
IT Institute, University of Salford, Salford, M5 4WT
R.S.Aylett@iti.salford.ac.uk

Abstract

We briefly consider thae nature of a social agent, summarise work carried out at Salford, and consider the communication requirements for different levels of socila organisation. We examine the ngineering implications for robot experiemnts of these communication requirements.

What are social agents?

We consider social agents in general to be those involved in interaction between same species individuals for extended periods in common locations. In nature social interaction is mediated by life events and tasks: mating and reproduction together with subsequent care of young, food gathering, protection against predators including sometimes construction of protective artefacts. But social interaction may not just be a question of external tasks: the social domain, once existing, may develop its own internal structure and relationships often resulting in a social hierarchy expressed and mediated through particular social behaviours and rituals. Social intelligence is then the ability of the individual to thrive in a social environment, both from the point of view of life tasks (e.g. successful reproduction) and within the social domain itself (e.g. high or rising social status).

When we consider artificial agents, whether entirely software based or also embodied in hardware as with robots, we may continue to approach social interaction from a agent-centred perspective as described above. However there is now an alternative point of view: the integration of such agents into a specifically human social milieu. While agent-centred perspectives are usually concerned with theoretical investigation, the human-centred perspective may also have a pragmatic focus in terms of meeting specific human purposes. Moreover the 'interaction between same-species individuals' basic to the agent-centred perspective may be completely absent, with the productive analogy being less natural animal or insect societies and more pets or other domesticated animals.

Social Agents at Salford

Work has been carried out for a number of years at Salford in the field of cooperative task execution with multiple robots (Aylett et al 97, Barnes et al 97) in which robots carry out cooperative object relocation tasks within a multi-agent architecture incorporating behavioural architectures on the robots and a fixed predictive planner agent supervising the task at an abstract level. This type of multi-robot system is aimed at industrial environments, but what is often forgotten here is that these environments are already complex social organisations in which the cooperating robots must also cooperate with human workers. Thus this work correctly seen falls into a human-centred approach rather than a purely agent-centred one.

Secondly, in the last year the author has worked with colleagues at VUB in Brussels on the origins of language and the grounding of adaptive generation of vocabulary in real world robots (Steels & Vogt 97). This is very much an agent-centred approach in which communicative competence between same-species individuals is the focus.

In both cases, communication is a major issue, whether between agents alone or between agents and humans. If social expertise consists in 'doing the right thing', then an important part of this is communicative activity.

Underlying this work are two premises, or perhaps more correctly, hypotheses. Firstly that social organisation is supported, developed and expressed by a communicative spectrum all the way from stigmergy (Holland & Becker 96) to human language. At each point on this spectrum there are different requirements in terms of the functionality required of the individual as well as different features in terms of the social expertise supported and expressed. We will briefly discuss these relationships.

The second hypothesis is that the behavioural substrate we think of as lying at the more primitive end of this spectrum is essential to the conscious activity we locate at the advanced end, not only historically but also functionally. This hypothesis has a number of practical consequences. One is that embodiment in the real world is vital to experimental activity - in spite of the difficulties involved in this - and that simulations may on occasion be misleading. Another is that making a service robot 'sociable' requires more than merely giving it a formalised body of social knowledge at a symbolic level.

It sometimes appears that workers in the field identify social expertise with cooperative task execution. A functional analysis of what a social organisation is 'good for' may support this approach, and at an intuitive level it is clear that 'many hands make light work' has been a powerful force in human social organisation. A clear functional role can also be identified for communication systems, from the alarm call of a blackbird to the planning of a family day out.

Yet it is important to remember as argued above that social expertise may also be reflexive, that is aimed at the social domain itself. Grooming in primate societies, it is argued (de Waal 82), has less to do with personal hygiene than with status, social hierarchy and social differentiation. Indeed it has been argued (Dunbar 93, Aiello & Dunbar 93) that the origins of language and of consciousness lie in the social domain itself rather than in cooperative task execution, that language evolved as a more efficient form of grooming and that consciousness was based on the evolutionary advantage in a social group of being able to 'put yourself in someone else's shoes'.

The robot dimension:

What does the investigation of social expertise using real robots require of the real robots? At the bottom end of the communicative spectrum, it requires the ability to use the environment as a communicative mechanism. It is at this level that most existing work is being carried out since it is possible to produce and perceive stimergic effects with extremely primitive actuation and perception mechanisms. For example, the cooperating robots in the work at Salford are able to jointly carry an object using force feedback from the object rather like two people carrying a table (Barnes 96). They require no ability to differentiate each other from the rest of the environment, a factor one could argue is inherent in stigmergy as a mechanism. This lack of differentiation means that hierarchy in this sort of social organisation is almost entirely functionally determined: if a bee is a drone then it cannot be a queen.

Even at this level, modes of interaction with the environment are usually fairly primitive, particularly if robots are compared to social insects for example. Robot sensory and actuation capabilities mean that they relate to the environment in an inflexible and impoverished way such that one might seriously question what sort of 'embodiment' is actually achieved. For example very few robots have any compliance in their effectors (an area mainly studied in fixed manipulators with some examples in multi-legged mobility), which reduces touch sensing down to the binary values generated by bumpers. The use of smell would require cheap but reasonably sensitive chemical sensors as well as the incorporation of generation mechanisms in robots. This is an area that seems hardly investigated at all. Sound is also a modality that is rarely used.

Problems of cost, power limitations, and the need to conserve a limited stock of robots rather carefully, rather than going in for the large-scale redundancy of individuals seen in insect societies also have the effect of constraining the type of interaction with the environment available. A robot society which could construct artefacts of the complexity of termite hills is some way off, with flocking and herding along with wandering and box-pushing the most common behaviours demonstrated using stigmergy. Many experimental systems lack any effectors beyond (usually) wheels or (less often, legs). From an engineering perspective, we have yet to construct robots anything like as physically sophisticated as ants, slugs or flies.

Species awareness

Further along the communicative spectrum robots require the ability to separate out 'others like me' from the rest of the environment, in other words species or perhaps merely social group awareness. An alarm signal, for example, is an explicit communication which only has meaning to others of the same species or social group. Here there are further engineering implications for robots since species recognition requires both characteristic features and the ability to recognise them. A very simple implementation might be to use the range-finding sensors available on most robots for obstacle avoidance. Thus, in a simple artificial environment, other robots may be the only sources of infra-red or ultra-sound emission.

Unfortunately, the very use of these output devices for obstacle avoidance means that the perceiving robot has to stop and turn off its own emitters in order to make a distinction between new emissions and the return of its own signals (Steels & Vogt 97). This seems very restrictive and in a sense misses the point: the difference between this mode of communication and stigmergy is precisely that species or group members can be distinguished from the environment, not that they can only be sensed if the environment is ignored.

A further option is vision (Belpaeme & Birk 97), but here the difficulty lies in the demands it makes upon on-robot processing: the amount of data delivered from a vision system and the computational demands of the algorithms used tends to lead to unacceptable power demands and slow processing. On the other hand, connecting the robot to an external processor severely impacts its autonomy. Comparing the sensory equipment generally available on robots with that available to animals who use this type of communication it appears that investigation of the other sensory modes already mentioned - such as sound and smell for example - and the combination of small amounts of data from a number of different sensors may provide a better way forward than the resource-demanding use of vision.

Individual recognition

Moving still further across the communication spectrum we arrive at language, which requires the establishment of individuality (Dautenhahn 95), a yet stronger requirement. Equipping each robot with a unique frequency infra-red or ultra-sound transmitter is one possibility, but it also seems likely that solving the species recognition problem through sensor fusion as suggested might lay the basis for individual recognition too. True, one can always cut through the problem by incorporating hard-wired symbolic names and a radio link, but arguably one might just as well use a simulator as incorporate a deus ex machina like this on real robots.

The development of individuality implies an extended 'life' and accumulation of experience (Dautenhahn 95) as well as an internal 'body image' - a representation within the robot of its own material form. Again, this has engineering implications as well as computing ones: an extended life requires the ability to recharge batteries unless autonomy is vitiated with an off-robot source of power. It probably also requires the use of static memory to retain data through periods of power drain. This is aside from the challenge of deciding what should be learned using what algorithms and what representations. The development of a 'body image' also has sensor implications - the measurement of battery level is the only internal sensor on mobile robots in most cases with nothing to give data on the spatial position of body parts except for rare cases of the use of inertial sensors to measure orientation (since the majority of mobile robot platforms run on level surfaces using wheels this must usually appear superfluous).

Conclusion

What this piece has tried to argue is that if one regards embodiment as important to communication and thus to social organisation, it raises a number of significant questions about the required functionality of robots to be used in experiments. In practice these are difficult questions: it is difficult to answer them without substantial expense and engineering backup. We argue that difficult issues must be confronted at quite primitive levels of interaction before one can provide the basis for higher level social interaction. It is perhaps not surprising that so much work is carried out in simulation, but it remains to be seen how far the conclusions of simulated work will be supported eventually on real-world robots.

References

Aiello, L. & Dunbar, R.I.M. (1993) Neocortex size, group size and the evoluion of language. Current Anthropology, 34, pp184-93

Aylett, R.S; Coddington, A.M; Barnes, D.P. & Ghanea-hercock, R.A. (1997) Supervising multiple cooperating mobile robots. Proceedings, Autonomous Agents '97, ACM Press, pp 514-5, 1997

Barnes, D.P. (1996) A behaviour synthesis architecture for cooperant mobile robots. Advanced Robotics and Intelligetn Machines, eds J.O.Gray & D.G.Caldwell, IEE Control Engineering Series 51.295-314 1996

Barnes, D.P; Ghanea-Hercock, R.A; Aylett, R.S. & Coddington, A.M. (1997) 'Many hands make light work? An investigation into behaviourally controlled coooperant autonomous mobile robots'. Proceedings, Autonomous Agents '97, ACM Press, pp 413-20, 1997

Belpaeme, T. & Birks, A. (1997) A real-world ecosystem featuring several robot species. Proceedings, European Conference on Artificial Life, ECAL '97, Brighton, July 1997

Dautenhahn, K. (1995) Getting to know each other - Artificial social intelligence for autonomous robots. Robotics and Autonomous Systems 16 (1995) pp333-56

Dunbar, R.I.M. (1993) Coevolution of neocortical size, group size and language in humans. Behavioural & Brain Sciences, 16 pp681-735

Franklin, S. (1996) Coordination without communication. Proceedings, The First UK Workshop on Foundations of Multi-Agents Systems, FOMAS '96, pp8-13, University of Waewick, October 1996

Holland, O. & Beckers, R. (1996) The varieties of stigmergy: indirect interaction as a control structure for collective robotics. Robotica - to appear

Steels, L. & Vogt, P. (1997) Grounding adaptive language games in robotic agents. Proceedings, Conference on Artificial Life, ECAL '97, Brighton, July 1997

de Waal (1982) Chimpanzee Politics: Power and Sex among Apes. London: Jonathan Cape