A "subsumption" architecture meeting the above requirement (2) has been proposed in the article "A Robust Layered Control System For A Mobile Robot" by R. A. Brooks, IEEE Journal of Robotics and Automation, RA-2(1), pages 14-23, 1986. Moreover, an architecture meeting the above requirements (1) and (2) has been proposed by the present inventor in "Phase Transitions in Instigated Collective Decision Making", in "Adaptive Behaviour", 3(2), 1995, pages 185-223. However, in general it is not possible to pre-program a set of routines or responses which will be appropriate in all the circumstances which a machine or computer program will encounter in practice. Moreover, in the human and animal kingdom genetic information does not provide behavioural responses appropriate for all circumstances which will be encountered during life.

Observation shows that in the human and animal kingdom there is an additional factor, namely the ability to specialise the innate primitive responses based upon experience and, in particular, based on interactions with others. A typical example of this ability is seen in the behaviour of young vervet monkeys. Initially, young vervet monkeys produce an alarm call when they see any flying birds, including ones which are harmless. However, they quickly learn to emit the alarm call only when flying predators are seen, much too quickly to be explained by the young monkey's own experience. It has been postulated that the young monkeys learn to specialise their responses based on the responses of their older peers who ignore the "false alarms" (see "The ontogeny of vervet monkey alarm calling behaviour: A preliminary report" by R.M. Seyfarth & D.L. Cheney, in "z. Tierpsychology", 1980, 54, pages 37-56).

A third requirement for "social intelligence" or appropriate social behaviour can thus be postulated:

(3) the ability to specialise responses originally triggered by primitive responsive behaviour, through interaction with others, and to remember the behavioural pattern as a secondary response.

The creation of a secondary response involves conditioning based upon inputs received from the outside. The generation of secondary responses represents conditioning of a system. Now, the ability to develop secondary responses based on sensory inputs in an unsupervised learning process involves the learning agent in an attempt to correlate a number of sensory inputs with a number of internal structures in an attempt to extend the knowledge base of the system. However, the computations involved in this correlation process are complicated. A "focus of attention" method for unsupervised learning has been proposed in "Paying Attention to What's Important: Using Focus of Attention to Improve Unsupervised Learning" by L. N. Foner and P. Maes, in "From animals to animats" 3, pages 256-265, ed. D. Cliff, P. Husbands, J. A. Meyer, and S.W. Wilson, 1994, the MIT Press. This method seeks to make the correlation task manageable by focusing attention on a limited number of factors. The proposed method is based on cognitive selectivity and employs world-dependent, goal-independent and domain-independent strategies.

An associative control process (ACP) for conditioning has been proposed in "Modelling Nervous System Function with a Hierarchical Network of Control Systems that Learn" by Klopf et al, in "From animals to animats" 2, pages 254-261, ed. J-A Meyer, H.L. Roitblat and S.W. Wilson, 1993, the MIT Press). The proposed ACP network includes two kinds of learning mechanisms, drive-reinforcement learning (in reinforcement centres) and motor learning (in motor centres). However, the proposed system lacks an internal driving force inciting the network to undergo conditioning and has no mechanism for focusing attention whereby to reduce the complexity of the correlation processes inherent in conditioning.

Another conditioning system is proposed in "No Bad Dogs: Ethological Lessons for Learning in Hamsterdam" by B. M. Todd and P. Maes, in "From animals to animats" 4, pages 295-304, ed. P. Maes, M. Mataric, J. Meyer, J. Pollack and S.W. Wilson, 1996, the MIT Press. This method involves the use of pre-defined "Behaviours" which are arranged in Behaviour Groups so as to be mutually inhibiting. The Behaviour Groups in their turn are arranged in a loose hierarchical fashion. The resultant structure is very complex.

The article "Reinforcement learning: A Survey" by L.P. Kaelbling, et al in the Journal of Artificial Intelligence Research", 1996, 4, pages 237-285, discusses a process whereby an agent can learn behaviour through trial-and-error interactions with a dynamic environment. This is a type of unsupervised learning. In each of a succession of time periods, the agent receives a set of inputs as well as a reinforcement signal (a "reward") and chooses an action. If the action is successful then the "reward" is allotted equally to all of the units which contributed to the choice of the action. The agent seeks to choose the action which increases the long-term sum of values of the reward, that is, of the reinforcement signal. It will be seen that, according to this proposal, the system conditions itself (adapts its responses to external stimuli) without external motivation, seeking merely to maximise an internal measure of the "success" of its action. However, such an approach does not take into account whether or not the conditioned response is "appropriate" or "successful" from the point of view of external agents with which the system interacts.

In the opinion of the present inventor, an agent emulating "social intelligence" not only should meet the above requirements (1) and (2) but also should meet a modified version of the third requirement, stated as follows:

(3') the ability to specialise responses originally triggered by primitive responsive behaviour, by reacting to the presence or absence of some <u>expected inputs</u> from the outside (typically from others), and to remember the behavioural