Summary of research record

Alan H. Bond

My research started with a PhD in theoretical physics at Imperial College, London University (I was a member of the elementary particle theory group lead by Abdus Salam and Steven Weinberg. They received the 1978 Nobel Prize for their work during this period), and continued with research in the computer science department at Carnegie-Mellon where I became acquainted with the ideas of Newell and Simon. This lead to an interest in developing a scientifically grounded artificial intelligence and psychology. My research history consists mainly of a long range search for a scientifically grounded computational model of the brain, embellished with a sequence of studies that I came across along the way. In addition, I have worked on some applied AI projects unrelated to this main goal.

My computer science research started at CMU and then continued at Queen Mary College in London, where I was tenured faculty in computer science. I tried to follow Newell and Simon by doing my own protocol experiments for problem solving and chess. One student reimplemented and extended GPS. Other students built autonomous robots. In 1978, I wrote a long position paper ‘‘An approach to artificial intelligence’’ attempting to ground AI in the notion of an autonomous surviving robot. This paper posited four bases - survival, real-time control, parallel architecture and learning. With a student, I was PI for a project to develop an intelligent rule-based learning robot, reported in IJCAI81. I worked with several parallel machines from 1978, and in 1983 designed my own SIMD architecture without a central control unit.

I collaborated with Les Gasser in editing the first book on multiagent systems, including the first in-depth review of multiagent concepts, which was published in 1988. I pursued two research ideas in multiagent systems, in order to develop a computational approach to social relationship. The first was to develop a notion of commitment among agents, published in 1990, and the second was to develop and implement a negotiation logic based on joint proof.

In 1990, I wrote another long position paper ‘‘What I have in mind’’ which discussed a comprehensive set of psychological ideas such as sequential processing at the top level, emotions as mental states, the representational needs of social interaction. This was an attempt to get closure on a creative synthesis. I had however gradually realized that I needed concepts and constraints from the hardware level in order to develop a computational model.

In 1992, I wrote down a parallel architecture that was inspired by blackboard ideas and also the modular architecture of the brain. I implemented this model in 1993 and based it on joint action with other agents. Since then I've extended the model to represent space, and to do problem solving. I've improved its efficiency to allow the development of applications. Its current cycle time and response time is about 100 milliseconds. I've started two kinds of application project - cooperation with virtual humans, and socially interactive robotics.

I have also worked on an industrial vision system which lead to fast grey-scale recognition of industrial parts, and also a hardware design concept, stream architecture , which computed by streaming video data through a series of frame stores with special purpose logic. At UCLA, I supervised an intelligent CAD group and we developed a complete CAD/CAM system in Prolog. This resulted in several papers. I also liaised with most of the Aerospace companies in LA. Lockheed used our system for one of their internal research projects. The other project was a data model for engineering databases, a collaboration with Charles Eastman in the UCLA architecture department. We developed a hierarchical data model with multiple dimensions and multiple inheritance, which was implemented. I also looked at parallel CAD models, using parallel logic programming on a hypercube and a shared memory machine. I also developed a computational approach to modeling design and manufacturing organizations, an exploratory paper was published in the DAI workshop 91.