

limitations of artificial intelligence, and these are likely to hinder progress in the design of automatic systems until well into the next century.

Artificial intelligence requires a two-step process: the representation of a domain of knowledge in computational form, followed by the design and implementation of algorithms to solve problems. Thompson sees the way in which decision-making is based on knowledge as imposing a crucial limitation on the further development of artificial intelligence.

Basically he sees knowledge as lying on a continuum from specialist to causal. Data at the specialist, structural, compositional or evidential end of the spectrum has been successfully utilised by artificial intelligence. Expert systems in word processing, mineral geology, bacterial infection diagnosis and computer installation are examples of artificial intelligence applications using specialist knowledge. In each of these systems, knowledge is represented in system specifications by "if-then" rules, and is manipulated by a variety of mechanized formal inference procedures. The other end of the spectrum, however, is general, functional, contextual, causal, situated or "common sense" knowledge, which entails understanding of use or significance. As yet the only decision-making system which uses general knowledge is the human being; artificial intelligence methodology has been unable to do so. Researchers do not know how to encapsulate the world in the machine. Of course there is no absolute division between the two forms of knowledge, but the further artificial intelligence systems diverge from the specialist end of the spectrum, the less successful are the results.

The artificial intelligence systems described thus far can be used for decision-making only in an advisory capacity. Since they utilise only the narrow specialist type of knowledge, they would not be good candidates for automatic launch-on-warning. Nonetheless, proposals to use them in that way are currently under consideration. Thompson was concerned that automatic launch-on-warning would empower inadequate machines to make the most crucial decisions facing humanity. He believed that what was essentially a political problem should not be delegated to technology. Flawed as humans are, they are far more capable of making sensible decisions than the best machine.

Joseph Weizenbaum suggested in his presentation that Ornstein and Thompson, if anything, probably erred on the side of overestimating the reliability of computers. He also pointed out that the unreliability of computers could cause unpredictable consequences; the recent failure of the computer system of a small US bank had had disastrous effects on the precious metals market.

He observed that Ornstein and Thompson had failed to note that large computer systems are incomprehensible in their entirety, and that no amount of study, simulation, or missile conferencing can overcome this. Larger computers are not so much deliberately designed as evolved, and