SSB: Fluid Concepts and Creative Analogies, Chapter 4

Jacob Peck

November 16, 2011

page 169 One of the deepest problems in cognitive science is that of understanding how people make sense of the vast amount of raw data constantly bombarding them from their environment.

page 170 The study of high-level perception leads us directly to the problem of mental *representation*. Representations are the fruits of perception.

page 171 *Perception may be influenced by belief.*

page 172 The distinguishing mark of high-level perception is that it is semantic: it involves drawing *meaning* out of situations.

page 173 How is it decided which subsets of the vast amounts of data from the environment get used in various parts of the representational structure? Naturally, much of the information content at the lowest level will be quite irrelevant at the highest representational level.

page 174 The formation of appropriate representations lies at the heart of human high-level cognitive abilities. It might even be said that the problem of high-level perception forms the central task facing the artificial-intelligence community: the task of understanding how to draw *meaning* out of the world.

page 175 The Physical Symbol System Hypothesis (Newell & Simon, 1976), upon which most of the traditional AI enterprise has been built, posits that thinking occurs through the manipulation of symbolic representations, which are composed of atomic symbolic primitives.

page 176 We are deeply skeptical, however, about the feasibility of such a separation of perception from the rest of cognition.

page 177 When BACON performed its derivation of Kepler's third law, the program was given only data about the planets' average distances from the sun and their periods. These are *precisely the data required to derive the law*.

page 178 It is difficult to believe that Kepler would have taken thirteen years to make his discovery if his working data had consisted entirely of a list where each entry said "Planet X: mean distance from sun y, period z". If he had further been told "Find a polynomial equation relating these entities", then it might have taken him a few hours.

page 179 The case of BACON is by no means isolated — it is typical of much work in AI, which often fails to appreciate the importance of the representation-building stage.

page 180 In cases such as these, it seems that no single, rigid representation can capture what is going on in our heads.

page 181 The mapping process, in contrast, is an important object of study especially because of the immediate and natural use it provides for the products of perception.

page 182 Even for a computer program, the extraction of such common structure is relatively straightforward.

page 183 Imagine what it would take to devise a representation of the solar system or an atom independent of any context provided by the particular problem.

page 184 Thus, when one is designing a representation for SME, a large number of somewhat arbitrary choices have to be made. The performance of the program is highly sensitive to each of these choices.

page 185 Such criticisms apply equally to most other work in the modeling of analogy. It is interesting to note that one of the earliest computational models of analogy, Evans' ANALOGY (Evans, 1968), attempted to build its own representations, even if it did so in a fairly rigid manner.

page 186 In cognitive science and elsewhere, scientists usually study what seems within their grasp, leaving problems that seem too difficult for later.

page 187 The only solution would be for the representation module to always provide a representation all-encompassing enough to take in *every possible aspect* of a situation.

page 188 The problem is simply that a vast oversupply of information would be available in such a representation. To determine precisely which pieces of that information were relevant would require a complex process of filtering and organizing the data available from the representation. *This process would in fact be tantamount to high-level perception all over again.*

page 189 One might thus envisage a system in which representations are gradually built up as the various pressures evoked by a given context manifest themselves.

page 190 This does not mean, however, that one must admit defeat. There is another route to the goal. The real world may be too complex, but if one *restricts the domain*, some understanding may be within our grasp.

page 191 On the one hand, like connectionist models, Copycat consists of many local, bottom-up, parallel processes from whose collective action higher-level understanding emerges. On the other hand, it shares with symbolic models the ability to deal with complex hierarchically-structured representations.

page 192 The Tabletop program takes a few steps further than Copycat does towards lower-level perception, in that it must make analogies between visual structures in a two-dimensional world, although this world is still highly idealized.

page 193 Research in artificial intelligence has often tried to model concepts while ignoring perception.