sities that are introduced by the specific context in which an object is encountered. Indeed, our empirical work tries to minimize such contextual considerations. Ultimately, a theory of object concepts will require a representation of objects in their ecological contexts, but our current theoretical goals are more limited.

CONCLUSIONS

Dependency structures can be hierarchically arranged in the sense that parts and aspects of concepts can have their own dependency structures. For instance, in the apple graph of Figure 1, two subnetworks of features can be discerned, one concerning the reproductive aspects of apples and the other containing the food-related features of apples. Each of these aspects of apples enforces its own coherence constraints which contribute to the overall coherence of the apple concept (cf. Rumelhart, Smolensky, McClelland, & Hinton, 1986).

As proponents of the view that conceptual coherence emerges from the interaction of multiple local (pairwise) dependency relations, we have no obvious way of distinguishing those relations that are intrinsic to a concept from those relations that are known but are not part of a concept's internal structure. For example, the relation between a ram's horn and the acoustical properties of a musical instrument do not seem intrinsic to the concept of ram and would not participate, we expect, in the determination of the mutability of the features of the concept corresponding to ram. This form of the binding problem is well-known and not easily solved by any learning device. Our belief, along with Thagard (1989), is that the dynamical properties of the kind of constraint satisfaction systems discussed by Hopfield (1982) and reviewed by Rumelhart and McClelland (1986) afford a useful perspective on this issue. If a concept is best conceived as an attractor in a high dimensional state space—as the solutions to Equation (1) are—then the set of relevant relations for a concept are jointly determined by the constraints imposed by all the relations in parallel.

The success of our abstract network representation provides more evidence for the utility of conceiving of concepts as attractors in a large-dimensional state space. One advantage of such a conception is that it illustrates how concepts can be both flexible and have structure. A concept emerges as multiple constraints are simultaneously satisfied. We have posited that the key constraints are pairwise dependencies between features. Concepts are flexible because little depends on some features, so they are easily transformed. Concepts, nevertheless, have structure because much depends on other features; they are relatively immutable.

Acknowledgments: This work was supported by a Richard B. Salomon Faculty research Award and by a grant from Brown University to Steven Sloman. Some of the results from Study 2 were reported at the Seventeenth Annual Conference of the Cognitive Science Society and the results from Study 4 at the Nineteenth Annual Conference. We would like to thank Nick Haslam, Doug Medin, Dorrit Billman, and two anonymous reveiwers for their comments on prior drafts.