Discovering Compositional Structure

by
Matthew T. Harrison
B.A., University of Virginia, 1998
Sc.M., Brown University, 2000

Doctoral dissertation
Ph.D. Advisor: Stuart Geman

Submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy
in the Division of Applied Mathematics at Brown University

Providence, Rhode Island
May 2005

Chapter 8

Conclusion

The general theme of this thesis was to suggest methods for discovering compositional structures in data. The unsupervised learning algorithms in the first part of the thesis hinted at some principles that might be useful for building compositional priors for image analysis. In particular, spatial dependencies were used to discover new compositions of old parts and temporal dependencies were used create invariant parts.

In future work, we hope to combine the two methods in order to create a truly compositional representation. This presumably involves the important and difficult task of designing computational principles for operating on a compositional hierarchy. We anticipate that creative solutions to the Markov dilemma, or the binding problem, will be crucial for both learning and computation within this framework.

Another line of future work will be to investigate information theoretic approaches that can unify selectivity and invariance. The principle of temporal stability offers one solution, but it is not clear how temporal stability formally relates to more statistical goals. Since, as we have shown, dependencies can be used to create both selectivity and invariance, it should also be possible to create an information theoretic criterion that can discover both.

The statistical methods in the second part of the thesis were designed for trying to understand how the brain might create compositional representations. Agnostic methods for investigating the receptive field properties of visual neurons are severely constrained by data limitations. We suggested that certain properties of natural images might facilitate their use. We also introduced a new class of jitter techniques which might be useful for investigating synchrony solutions to the binding problem.

In future work, we hope to experiment more with these methods on real data. The agnostic methods will undoubtedly need substantial work before they are practical. There may also be unanticipated physiological constraints other than mere sample size that severely limit their utility. The jitter methods, on the hand, are immediately applicable for spike train analysis. They should be useful for rate-controlled measures of synchrony regardless of how the brain actually uses synchrony, if it uses it at all.

Compositionality is an important, yet elusive principle for general pattern recognition. We believe that it might underlie the remarkably fast visual learning demonstrated by a variety of biological vision systems. Hopefully this thesis will contribute to increased interest in compositionality within the broader vision community.