Circuit Sharing and the Implementation of Intelligent Systems

Michael L. Anderson

Department of Psychology, Franklin & Marshall College, Lancaster, PA, and Institute for Advanced Computer Studies and Neuroscience and Cognitive Science Program, University of Maryland, College Park, MD 20742, USA
michael.anderson@fandm.edu

Abstract

A foundational and highly contested question in cognitive science is whether and to what degree the brain is a modular system. This talk outlines some of the broad architectural implications of the modularity thesis, and reports on an attempt to test for them. The results indicate that the same brain regions contribute to functions across various cognitive domains, but in each domain cooperate with one another in different patterns. This does not appear to be compatible with the modularity thesis. The talk discusses the implications of the finding for the best approach to analyzing, modeling and reproducing cognitive functions in intelligent systems.

Architecture of Modularity

Insofar as modules are domain specific and implemented in dedicated neural substrates (however distributed in the brain), this suggests they will be supported largely by non-overlapping elements. Differences in cognitive function result from differences in the neural circuitry involved in implementation. In contrast, redeployment suggests that neural circuits were frequently reused in evolution, whenever an emerging function could be realized more efficiently this way (Anderson, 2007). For redeployment, differences in function come primarily from differences in the patterns with which the same brain areas cooperate under different circumstances. Given this contrast one can ask simply: do the differences in the neural implementations of different behavioral domains appear to be more attributable to differences in the neural real-estate used, or to differences in the ways in which the same brain regions cooperate in each case?

Methods

The method is simple: choose a spatial segmentation of the cortex (current analyses use Brodmann areas); analyze a large number of fMRI studies to determine the baseline chance for coactivation of each pair of regions; for different levels of an independent variable of interest, use chi-squared to find pairs where chance of co-activation significantly differs from baseline; and represent the results as graphs, where nodes are the regions and edges indicate significant co-activation (Anderson et al., in press; *see also* Toro et al., in press). Here we analyzed 472 experiments in eight cognitive domains, and built co-activation graphs for each domain. We compared these graphs to see if their differences were primarily reflected by different nodes, or different edges.

Results and Discussion

Briefly put, the graphs shared over 80% of their nodes, but only 15% of their edges. This result favors the redeployment view: in the brain, cognitive functions are implemented in largely overlapping neural substrates. The result is also compatible with a PDP framework (e.g., Mesulam, 1990); however earlier results (Anderson, 2007) fit redeployment better than PDP. This implies that the best way to approach the understanding/modeling of cognitive function is *not* by investigating behavioral domains (e.g., perception, language, motor control) in isolation from one another, but by looking for points of overlap in neural implementations, and exploiting these to guide the analysis and decomposition of the functions in question. Likewise when one engineers intelligent artifacts: rather than build parts (parsers, vision systems, etc.) in isolation, teams should consider how to build integrated systems out of shared low-level components.

References

Anderson, M.L. (2007). Evolution of cognitive function via redeployment of brain areas. *The Neuroscientist*, 13(1): 13-21.

Anderson, M. L., Brumbaugh, J. & Suben, A. (in press). Investigating functional cooperation in the human cortex with simple graph-theoretic methods. In: W. A. Chaovalitwongse, P. Pardalos and P. Xanthopoulos (Eds.) *Computational Neuroscience* Springer Verlag.

Mesulam, M.-M. (1990). Large-scale neurocognitive networks and distributed processing for attention, language and memory. *Annals of Neurology* 28: 597-613.

Toro, R., Fox, P.T., & Paus, T. (in press). Functional coactivation map of the human brain. *Cerebral Cortex*.