

realizability with respect to TALoNS as well. And, as Anderson acknowledges, these constraints often get trumped by cultural and social constraints that also get preserved and conserved across time. So I want to push back a little against the picture of brain regions evolving particular dispositions that get combined in various ways like primary colors to collaborate on different tasks.

Rather than speak of brains as control systems with intrinsic dispositions, I want to say that brains are hubs in the graph theoretic sense, but they are living hubs in multilevel dynamical processes that are extended in space and time. The networks in question are physio-bio-cultural, and brains are truly historical artifacts that bear the marks of their origins and development. In this model, extended organisms or agents spanning brain, body, and environment are the primary locus of control, and the structural elements that instantiate such networks are often secondary to the networks themselves (Silberstein & Chemero 2011).

This brings me directly to my second main point about the autonomy of folk psychology. Rather than reduction or elimination, neural reuse and extended cognition so conceived actually support functionalism in the sense that we individuate processes with regard to their role and effects as opposed to their structural constituents. Again, looked at in this way, we find a great deal of multiple realizability within and across human and non-human brains. Anderson places a lot of emphasis on the importance of selection in understanding brain function (p. 296), and I agree, but I think there are others aspects of evolutionary and developmental biology that are equally important and that do not fully reduce to selection. Evolutionary theory has its own version of functionalism in the form of convergent or parallel evolution such as mimicry and flying (McGhee 2011). Focusing on cognitive convergence, creatures with very different brains and selective histories seemingly converge on similar behavioral and cognitive strategies for dealing with “socioecological” problems. For example, corvids do not even possess a prefrontal cortex, yet they exhibit very intelligent behavioral and cognitive strategies similar to primates. There is growing evidence across the board that creatures with very different brains have in many respects converged on relatively similar minds (McGhee 2011). Well-known examples of this from insects to mammals include: tool use, architectural behavior, agricultural behavior, social or collective behavior, mathematical behavior, and language use (McGhee 2011, Ch. 6). There is also a growing consensus that a wide variety of different species with very different brain structures and nervous systems possess some form of not only sophisticated convergent cognition, but also consciousness, self-awareness, and metacognition. Examples abound, such as the mourning behavior of gorillas and dolphins (McGhee 2011, p. 240). I think the best explanation for cognitive convergence is that the affordances, environmental and social features often trump structural neural constraints whether imposed by physics or selection.

So, although I agree with Anderson that the brain must have its voice, what I think it is telling us is that neural reuse is best viewed as a subset of what developmental and evolutionary biologists call plasticity and robustness (Bateson & Gluckman 2011). Although there are many different kinds of both robustness and plasticity, in general robustness refers to relative stability or invariance across environmental, genetic, or cellular transformations, and plasticity refers to cases wherein features of the organism are held constant such as genotype, and yet because of environmental transformations the organism nonetheless manifests very different or unique adaptive traits or behaviors (Bateson & Gluckman 2011, p. 8). Different forms of plasticity include: phenotypic, molecular, variety of neural types, immunological, and behavioral (Bateson & Gluckman 2011, p. 46). Both of these features of complex biological systems are of course at the heart of the epigenetic revolution in biology. As many people have pointed out, robustness and plasticity are two sides of the same coin: “Plasticity is often regulated by robust mechanisms and robustness is often generated by plastic mechanisms” (Bateson & Gluckman 2011, p. 46). It is important to note that, although many biological mechanisms possess the

properties of robustness and plasticity, these features cannot in principle be explained mechanistically in terms of localization and decomposition. These are global/systemic and scale-invariant features of such biological systems. Convergent evolution, robustness, and plasticity all go hand-in-hand, and they all point to the strongly extended nature of phenotype, behavior, and cognition. For one take on what such a Gibsonian cognitive neuroscience might look like, see Silberstein and Chemero (2012).

Modularity in network neuroscience and neural reuse

doi:10.1017/S0140525X15001673, e133

Matthew L. Stanley^a and Felipe De Brigard^{a,b,c}

^aCenter for Cognitive Neuroscience, Duke University, Durham, NC 27708;

^bDuke Institute for Brain Sciences, Duke University, Durham, NC 27708;

^cDepartment of Philosophy, Duke University, Durham, NC 27708.

mls103@duke.edu felipe.debrigard@duke.edu

www.felipedebbrigard.com

Abstract: Neural reuse allegedly stands in stark contrast against a modular view of the brain. However, the development of unique modularity algorithms in network science has provided the means to identify functionally cooperating, specialized subsystems in a way that remains consistent with the neural reuse view and offers a set of rigorous tools to fully engage in Anderson’s (2014) research program.

In *After Phrenology: Neural Reuse and the Interactive Brain*, Anderson (2014) offers a valuable framework for understanding neural (re)organization and its relationship to cognitive functioning. This framework, neural reuse, maintains that most regions of the brain are involved in multiple and diverse cognitive tasks, and that the brain is capable of achieving functional flexibility precisely because it can redeploy the same neural structures for a variety of purposes (p. 5). This neurofunctional architecture allegedly stands in stark contrast against a modular view of the brain. However, in our view, recently developed, data-driven modularity algorithms in network science (modularity_{net}), which make use of mathematical formalisms from graph theory, remain consistent with Anderson’s view while providing a set of rigorous methods to engage in Anderson’s research program.

Modularity_{net} algorithms are computed on networks, which are formally described using graph theory (Newman 2006). All networks are composed of differentiable elements of the system (nodes) and pairwise relationships between those elements (edges). In the context of human functional brain networks, each node represents a discrete parcellation of brain tissue, and edges represent measured functional interactions (connectivity) between pairs of nodes (Bullmore & Sporns 2009; Rubinov & Sporns 2010). Modularity_{net} algorithms can be used to identify nodes that form tightly interconnected subgroups of interacting brain regions functionally cooperating to subservise certain tasks. These modularity_{net} algorithms provide valuable information about (1) the extent to which the entire system can be (partially) decomposed into modules_{net}, (2) the extent to which the nodes within each module_{net} are preferentially cooperating with nodes within their own module_{net} versus nodes in other modules_{net} in the system, and (3) the different functional roles of specific nodes within each module_{net} (Stanley et al. 2014; Telesford et al. 2011). As such, modularity_{net} constitutes an alternative to dimensionality reduction approaches discussed by Anderson (e.g., diversity variability, Dice’s coefficient), while providing more information about the properties of these functionally cooperating groups of regions subserving different tasks. Although modules_{net} are not isolated, autonomous, encapsulated processing units – as demanded by more traditional formulations of “modularity” – they do still perform particular, specialized functions during




Figure 1 (Stanley & De Brigard). Provides an illustration of possible properties of a network's modular_{net} organization. Suppose each individual node (circle) represents some discrete, predefined portion of the brain, and the links between nodes represent functional interactions between nodes. Node color represents the module_{net} to which that node belongs. Notice that the nodes composing the blue module_{net} are very densely interconnected among themselves, suggesting that those nodes are cooperating to achieve some function, despite remaining interconnected to all other nodes in the system via direct and/or indirect connections. In contrast, the nodes composing the red, orange, and green modules_{net} are not nearly as densely interconnected or clearly defined. The kind of modular structure exhibited by the group of blue nodes allows researchers to maintain the position that modules exist and are responsible for specialized, segregated functions while remaining consistent with Anderson's overarching framework that militates against purely decomposable, strictly domain-specific, encapsulated, and isolated modules.

certain tasks via the interactions between brain regions within each module_{net} itself. Specific modules_{net} are thought to perform specific functions, even though each module_{net} remains interconnected, or integrated, with all other modules_{net} in the system by a complex set of inter-module_{net} connections (Fig. 1). Thus understood, modularity_{net} safeguards a conception of specialized, segregated functioning, which is central to more traditional views of modularity. Indeed, it has been shown that a modular_{net} architecture both exists in functional brain networks and is advantageous because it increases the robustness, flexibility, and stability of the system (Barabasi & Oltvai 2004; Valencia et al. 2009).

In what follows, we briefly identify the fundamental principles of neural reuse and show that modularity_{net} is consistent with it. We also suggest that modularity_{net} provides the best set of methods for investigating neural reuse, while preserving attractive features traditionally ascribed to a modular view of the brain.

The *first* claim Anderson makes about neural reuse provides the foundation for his framework: *Brain regions should be used and reused for diverse purposes in various task domains* (p. 9). In using modularity_{net} algorithms to investigate neural reuse, we assume that (1) subgroups of nodes (brain regions) identified as highly interconnected during a task consistently across persons are being used for that task, and (2) nodes can change in module_{net} allegiance between any two conditions. If, in fact, nodes change in their functional properties to be part of different interconnected, functionally cooperating groupings (modules_{net}) from task A to task B, then the modularity_{net} algorithm will be able to identify this shift. Because modularity_{net} algorithms are data-driven, the quantity of nodes forming any given module_{net}, the spatial locations of nodes within any given module_{net}, the consistency of modular_{net} organization across a set of persons, and the extent of interconnectedness of nodes comprising each module_{net} are determined by the very nature of the functional brain network,

and not by the experimenter. Furthermore, modularity_{net} algorithms are designed to admit the possibility that there are no clear, distinct modules_{net} in the network. Consequently, modularity_{net} serves as both a validation of neural reuse between different tasks and as a way to identify those functionally cooperating groups of brain regions subserving any task of interest. Modularity_{net} makes no a priori assumptions about the selectivity of local neural structures for specific task(s). If, in fact, local neural structures are not highly selective and typically contribute to multiple tasks across domain boundaries, then modularity_{net} will show how neural structures (nodes) change modular allegiance by reorganizing their connections for any set of tasks.

Recent work has demonstrated that human functional brain networks exhibit a modular_{net} architecture, but that architecture is neither temporally static nor anatomically fixed. Instead, nodes alter their interactions with other nodes in the network to form new modules_{net} depending upon the demands on the system (Bassett et al. 2011; Cao et al. 2014; Meunier et al. 2014; Moussa et al. 2011; 2014; Stanley et al. 2014). For example, Stanley et al. (2014) demonstrated that the nodes within the module_{net} mostly composed of default-mode brain regions maintained a highly consistent, densely interconnected, functionally cooperating module_{net} during a minimally demanding working memory task. However, when more attentional and working memory resources were required, the previously observed default-mode module_{net} dissolved. During this more difficult working memory task, a different set of brain regions, many of which have traditionally been considered to be involved in working memory processes, became densely interconnected, forming a new, consistent module_{net} across persons. Changes in modular_{net} organization within individual persons have been shown to facilitate behavioral adaptation during simple tasks, further emphasizing the constantly changing, dynamic nature of modules_{net} (Bassett et al. 2011).

The *second* claim Anderson makes about neural reuse concerns the importance of interactions between different neural elements: *Functional differences between task domains are critically reflected in the different patterns of interaction between many of the same elements forming the right neural team for a job* (pp. 9, 46). Modularity_{net} is computed by detecting certain patterns in the interactions between all brain regions comprising the network simultaneously. That is, modularity_{net} takes into account all interactions between each and every brain region and subsequently identifies the subsets of brain regions that tend to be more densely interconnected (i.e., interacting more strongly) among themselves than the rest of the network. The functional groupings of highly interconnected brain regions observed during task A that appear consistently across subjects are thought to actually subserve the cognitive process(es) associated with task A. Although modularity_{net} does still provide relevant information about the spatial locations of nodes comprising each module_{net} in the brain, modularity_{net} provides more important information regarding how certain regions are functionally cooperating during any given task, the extent to which regions are densely interconnected, how that interconnectedness changes across tasks, and the relative importance of specific nodes within modules_{net} for facilitating integrated and segregated neural functioning.

Critically, however, if one grants that achieving a task is really about putting together the right "neural team" (i.e., module_{net}), then that module_{net} is engaged in specialized processing for that specific task. Different components of that neural team might be serving different subfunctions, but the components of the team are all working to serve the particular function carried out by the module_{net} itself. So, in investigating the neural basis of cognitive processes, one cannot merely be concerned with integration between modules_{net}, as Anderson suggests (p. 42). Modularity_{net} algorithms recognize the importance and mutually dependent necessity of both segregation (specialization) within a particular module_{net} and integration between modules_{net}. Understood in this way, segregation and integration of function are not wholly

separable. But, nodes in modules_{net} still densely interconnect to form functionally cooperating groups engaging in specialized functioning during specific tasks. By recognizing the importance and mutual necessity of both segregated and integrated functioning, modularity_{net} seems to preserve the one truly valuable tenet of more traditional views of modularity – that is, that there is, in fact, segregated, specialized functioning in the brain that is necessary for cognitive functioning.

The *third* claim Anderson makes about neural reuse implicates evolutionary assumptions: *Later emerging behaviors/abilities should be supported by more disparate neural structures* (p. 9). Modularity_{net} makes no a priori assumptions about which behaviors/abilities are supported by the most scattered set of structures in the brain. Importantly, modules_{net} need not be spatially contiguous in the brain, because edges in functional brain networks are defined as statistical dependencies in neural signal between nodes. Because the size, consistency, and spatial scatter of modules_{net} are determined by the nature of the functional network itself, modularity_{net} actually provides the ideal way to empirically test whether later emerging behaviors/abilities are supported by a highly interconnected, consistent module_{net} comprising a greater proportion of different structures broadly scattered throughout the brain. So, if Anderson's third claim is right, then the data-driven modularity_{net} algorithm should identify the functional groupings of nodes representing the latest developing modules_{net} (during the relevant tasks) as those that exhibit the most noncontiguous and densely interconnected, yet spatially consistent (across persons), spatial scatter throughout the brain.

The *fourth* claim Anderson makes maintains that neural reuse is a guiding functional principle across many different spatial scales: *Neural reuse does not go away, no matter how small the brain region* (pp. 29–30). By using modularity_{net} algorithms, it is possible to investigate neural reuse on many different spatial scales. Even among the existing network analyses of fMRI data alone, researchers have reported results from 70-node to 140,000-node whole brain networks (Stanley et al. 2013). Assuming adequate computational power, modularity_{net} can be computed on networks at any of these spatial scales, and different kinds of information can be extracted from the modular_{net} architecture (or lack thereof) observed at each scale. Recent work has also led to the development of hierarchical modularity_{net} algorithms (Arenas et al. 2008; Meunier et al. 2009; Sales-Pardo et al. 2007), where each module_{net} obtained at the partition of the highest level can further be decomposed into submodules_{net}, which in turn can be decomposed into subsubmodules_{net}, and so on. This makes it possible to rigorously compare modular_{net} organization at many different spatial scales in order to capture reuse at different levels of organization within the same brain during the same task.

The *fifth* critical claim Anderson makes about neural reuse concerns separate modifiability and decomposability: *Neural reuse does indeed militate against separate modifiability; the brain is not a nearly decomposable system consisting of separately modifiable parts* (pp. 39, 40). To understand the compatibility between Anderson's fifth claim and modularity_{net}, it is necessary to distinguish between strong and weak versions of decomposability. Decomposability_{strong} refers to a fully separable system in which any element can be removed or altered without significantly impacting the remaining elements of the system (e.g., a massive modularity view). Decomposability_{weak} refers to a system that has parts that tend to group together to perform specific functions, but in which each part remains connected to the rest of the system through a complex pattern of interactions, such that no functional group can be changed without changing the system as a whole. Although more traditional formulations of modularity operate under the assumption of decomposability_{strong} that further assumes separate modifiability, modularity_{net} only assumes decomposability_{weak}, which does not allow for separate modifiability. Functional brain networks are investigated as fully interconnected, interdependent, nonlinear systems. This means

that no module_{net} can be modified in practice without also changing the modular_{net} architecture of the system as a whole. By accepting decomposability_{weak} but not decomposability_{strong}, modularity_{net} is able to maintain a particular notion of segregated, specialized functioning while still remaining consistent with neural reuse.

Neural reuse holds tremendous promise as a framework with which to understand and investigate the neural bases of cognitive processes. We have argued that recent developments in network neuroscience offer a particular view of modularity – modularity_{net} – that is consistent with the fundamental tenets of neural reuse. Additionally, we have suggested that these modularity_{net} algorithms constitute an excellent tool for empirically investigating neural reuse, even for researchers reluctant to relinquish the idea of specialized modules altogether.

Neural reuse leads to associative connections between concrete (physical) and abstract (social) concepts and motives

doi:10.1017/S0140525X15001685, e134

Yimeng Wang and John A. Bargh

Department of Psychology, Yale University, New Haven, CT 06511.

yimeng.wang@yale.edu john.bargh@yale.edu

<http://www.yale.edu/acmelab/index.html>

Abstract: Consistent with neural reuse theory, empirical tests of the related “scaffolding” principle of abstract concept development show that higher-level concepts “reuse” and are built upon fundamental motives such as survival, safety, and consumption. This produces mutual influence between the two levels, with far-ranging impacts from consumer behavior to political attitudes.

The neural reuse theory has proven remarkably generative and supportive of research activity on embodied cognition, motivation, and behavior. We first present and discuss recent research evidence consistent with a “conceptual scaffolding” principle that is in harmony with the neural reuse perspective. We then discuss how conceptual development in neuroscience and empirical research in social psychology may collaborate in fathoming the neural structure of grounded cognition.

Anderson (2014) notes that conceptual metaphors – structure and logical protocols from one domain guiding and structuring those in another (e.g., life is a journey, love is war) – may be only one “side effect” (p. 26) of the neural reuse process on a larger scale. Evidence from social and consumer psychology on embodied forms of judgment and behavior suggests that higher-order mental conception is not built from the more basic functional models as “prototypes” (Anderson 2014, p. 17) as much as it is built upon (associatively connected to) the biological groundings of the primal functions, sharing their neural substrates by reuse. Several lines of research provide support for this proposition.

The conceptual scaffolding account (Williams et al. 2009), which was influenced by and is in harmony with the neural reuse principle, argues that more abstract concepts and complex, higher-order mental functions (person perception, self-esteem, value judgment) grow organically from and are hence “built upon” the more fundamental, innate needs, such as for survival, safety, and resource acquisition (consumption). The “built upon” notion leads directly to the assumption that associative connections will be formed between the concrete and the abstract concepts or goals, associations that remain intact throughout the lifespan. This in turn leads to hypotheses regarding the use of physical level concepts in describing more abstract social and psychological phenomena, as documented by extant theory and research on metaphor use (Lakoff & Johnson 1980; Landau et al. 2010). Social relationships (a *close* relationship, a *warm* father, a

