

Neural Reuse:
A New Perspective on the Fusiform Face Area

by

Josh Bowie

A thesis submitted to the Faculty of Graduate and Postdoctoral Affairs in
partial fulfillment of the requirements for the degree of

Master

of

Cognitive Science

Carleton University
Ottawa, Ontario

©2015
Josh Bowie

Table of Contents

Abstract.....	iii
Acknowledgements.....	iv
Introduction.....	1
Chapter 1 – Modularity.....	5
<i>Section 1: What is Modularity?</i>	5
<i>Section 2: Functional Versus Anatomical Modularity</i>	6
<i>Section 3: Fodorian Modularity</i>	8
<i>Section 4: Massive Modularity</i>	15
<i>Section 5: Evolution, Development and Learning: Support and Pitfalls</i>	22
<i>Section 6: A Critique of Modularity</i>	34
Chapter 2 – Neural Reuse	41
<i>Section 7: The Mental Mechanism – A Precursor to Neural Reuse</i>	44
<i>Section 8: General Theories of Neural Reuse</i>	48
<i>Section 9: Neural Reuse Framed by Homology and Development</i>	58
<i>Section 10: Are Workings Fixed?</i>	65
<i>Section 11: Granularity</i>	74
<i>Section 12: At Odds With Massive Modularity</i>	79
Chapter 3 – A Neural Reuse Perspective of the FFA: The Search for a Working	82
<i>Section 13: Fusiform Face Area</i>	85
<i>Section 14: The Role of the FFA: Competing hypotheses (A Within-Modality Analysis – Single Domain versus Multi-Domain)</i>	89
<i>Section 15: Similar Domain General Uses Across Sensory Modalities (A Between Modalities Analysis)</i>	95
<i>Section 16: Between-Domains Analysis and Connectivity of the FFA</i>	99
Conclusion and Next Steps	103
References.....	108

Abstract

This thesis examines the effectiveness of Fodorian modularity and massive modularity in explaining the brain's cognitive architecture, and compares these perspectives to that of neural reuse. Each is assessed against developmental and evolutionary principles, and against evidence from philosophy, psychology, neuropsychology and neuroscience. After concluding that neural reuse is better for understanding cognitive function than modularity, I carry out a case study of the fusiform face area (FFA) in an attempt to learn more about this brain area by putting the neural reuse perspective into practice. The findings suggest that the claim the FFA is a domain specific module for face recognition is not supported by the evidence, rather it is better described as a region that is specialized for disambiguation of within category information (e.g., faces, cars, animals). There is also some evidence that the FFA contributes in this way to functions beyond the visual modality (e.g., auditory perception).

Acknowledgements

I would first like to thank Andy Brook and Vince Bergeron for their infinite patience and guidance. My work benefited greatly from Andy's limitless knowledge and attention to detail, and from Vince's expertise on the topic of neural reuse and frequent encouragement. Andy and Vince have helped me to become a better thinker; I could not ask for more than this and I am truly grateful. I must also thank Marg, Bruce and Pauline for showing an interest in my studies. Regardless of whether it was genuine or not, it helped me get through the past three years. I'm also indebted to my uncle Peter, who sparked my interest in cognitive science and directed me toward Andy. Thanks to the staff at Starbucks coffee for being awesome. Thank you to Liane Dubreuil for holding my hand – without you, my lack of organizational skills would have certainly led to my early exit from the program. Finally, I would like to express my immense appreciation to Laura, who despite the burden placed on her during the birth of our two children (and in between) found it within her heart to support and celebrate my work, even during times when it took me away from changing diapers and other fatherly duties. And thank you to Noah and Darwin for being my inspiration and motivation.

Introduction

The turn from behaviorism to cognitive science in the late 1970s marked a new focus on the workings of the brain; a new and ardent effort to describe and explain a limitless list of cognitive functions. As research in cognitive science, neuroscience and neuropsychology advanced, it became clear that certain regions of the brain, particularly within the cortex, are linked to specific higher-level functions.

General functions (e.g., memory, emotion, language) have been attributed to coarse-grained anatomical sites, while more specific functions (performing tasks, perceiving faces, language production) are associated with finer grained neural networks. At some point, as we move from general functions to those that are more specific, the ability to easily recognize function becomes strained, if not lost altogether, as function can be so different from our understanding that it is unrecognizable. At times, this dismantling of the cognitive machine can present significant challenges. However, if overcome, these challenges can present profound new insights that not only change the rules of the game, but the game itself.

To elaborate further on this statement, by deconstructing functions into their component parts that provide some sub-function, we may find that these sub-functions explicate the higher-level function in a more precise way than by simply examining the higher-level function on its own – here we have a rule change, i.e., generally, language areas in the temporal cortex are responsible for language, but more specifically, there exists a localization of sub-functions (language production vs language understanding);

the latter clearly provides a deeper understanding. The game changes, however, when we explore further into the weeds. Take, for example, Broca's area, also known as Brodman area 44/45 of the frontal lobe, which is thought to be essential for language production. Research conducted by Patel (2003) and Schubotz & Fiebach (2006) has demonstrated that this area is likely responsible for a number of other functions across multiple domains, including syntactic operations in both music and language, object manipulation, action perception, etc. Suddenly, the game becomes less about language; in fact, language simply becomes one of many players on the field. So what is the new game? This question is difficult to answer. A best guess, as proposed by Shubotz & Fiebach (2006), is that the sub-function of this particular grouping of neurons is to process the hierarchical sequencing of inputs.

The corollary of this perspective on how the brain works is that it becomes natural to posit that at the sub-function level, opportunity for use of the function broadens. An analogy of a molecule might be useful. I will use the molecule glucose, however almost any molecule could have been chosen. The molecular formula of glucose is $C_6H_{12}O_6$, where C is carbon, H is hydrogen and O is oxygen. Water contains similar components, H and O, in the form of H_2O . While both H and O are contained within both molecules, they combine to form results that differ from each other in both their physical properties and the roles they play in the human body. The same can be said for sub-functions within the brain – they can be re-used to support multiple functions. The concept of neural reuse is a relatively new one in the field of cognitive science, however recent work by Anderson (2010 & 2014) has established a framework upon which further theorizing can expand our understanding of how the brain is organized and carries out certain functions,

while experimental research can provide the empirical evidence that is required to confirm or refute this concept.

While the field of chemistry and our interest in molecular composition and function is relatively old and advanced, the field of cognitive science is relatively young and immature. In cognitive science, our work is very much ahead of us. Defining the vast number of sub-functions and understanding the role they play in higher-level functions, I will suggest, is paramount to understanding how the mind works.

In this paper, I will attempt to expand our knowledge in this area by shedding new light on the sub-functions associated with the fusiform face area (FFA) and propose additional functions that may be provided at the whole system level, in particular with respect to expertise for objects and sensory input from non-visual modalities. I will take the perspective of neural reuse to examine a broad range of behaviours associated with neuroimaging results of the FFA, with the objective of inferring non-apparent functions of the FFA.

But first, in Chapter 1, I will dissect the model of cognitive architecture from which neural reuse has emerged – modularity. While modularity is not the only alternative perspectives used to describe cognitive architecture, it is best placed to reveal the strengths of neural reuse by comparison. It is my view that neural reuse falls along a temporal continuum, where, over time, we have expanded on earlier perspectives to develop more sophisticated explanations of how the mind works. Indeed, the measure along this continuum is ‘explanatory value’, which is measured in terms of how well a given perspective fits the facts, it’s ability to generate interesting predictions, and it’s propensity to generate interesting questions. I envision the left side of this continuum to

include Fodorian modularity, which was a first crack at explaining the architecture of the mind and described it as a doughy domain-general mass with discrete and impermeable modules that carried out a few specific functions, i.e., sensory and linguistic. The outcome of this perspective clearly defined the architecture of a small list of functions, but provided little explanatory value for how the rest of the mind works.

In the middle is massive modularity; an attempt to take the principles of Fodorian modularity and combine them with evolutionary psychology to create a perspective capable of describing hundreds of modules that have each evolved to respond to the environmental challenges of our ancient ancestors. Building on Fodor's thinking, this perspective comes with significantly more explanatory potential. However, what it adds in explanatory value is equaled by its ambiguity and propensity to confuse what it meant by a 'module', e.g., the introduction of 'wide scope encapsulation' (Carruthers, 2006), a weak type of encapsulation, which on the face of it looks like an oxymoron.

Finally, on the right side of this continuum is neural reuse. The neural reuse perspective has evolved from criticism of the massive modularity approach to explaining the mind. In particular, neural reuse advocates have identified problems with high-level fMRI results and the methodology followed by many researchers that assume a modular architecture in the design of experiments, which consequently leads to the reinforcement of this view. For example, due to the sensitivity of the instruments, the statistical measures used (multiple comparisons problem – when an apparently statistically significant observation occurs by chance due to a very large number of variables) and use of the subtractive method, some researchers are concerned that artifacts can sometimes be mistaken for meaningful activity (Bechtel, 2003). An example of this is a study in which

a dead salmon demonstrated meaningful fMRI activity in a perspective-taking task, highlighting that random noise can influence results if multiple comparisons are not controlled for (Bennett et al, 2009). This apparent neural activity was due to random false positive readings for a small number of voxels (16 of 8064). In addition, modularity perspectives (massive modularity in particular) often leads researchers to be content with assigning one function to a specific brain region, rather than continuing to search for other uses of the region. This bias can lead to less imaginative research projects and limit progress.

By casting these biases aside, and by altering the way in which we examine brain regions, the neural reuse perspective provides greater explanatory value for cognitive function than modularity perspectives. Given the importance I have ascribed to neural reuse's precursor perspectives, I will dedicate Chapter 1 to describing the first two stages on the explanatory continuum: Fodorian modularity and massive modularity.

Chapter 1 – Modularity

Section 1: What is Modularity?

Modularity has dominated the conversation of how elements of the mind are packaged within the brain, and how these discrete packages process and transfer information. Conceptually, it is advantageous to organize the brain into anatomical parts that are associated with functions that are both observable and measurable. However, precisely what the term “module” means has led to an ongoing controversy in the field of cognitive science. In particular, this tension lies in the permeability of the membrane (if I

may use this metaphor) that encloses a module. Here, permeability refers mainly to the information that is used by a module; in other words, its level of *encapsulation*.

Equally relevant in this vein is the extent to which a proposed module restricts its input to a particular type of stimuli. Indeed, the level of *domain specificity* is the primary marker by which many cognitive scientists and philosophers determine the likelihood that a given brain region is a module. In other words, a brain region is modular if it processes a distinct and narrowly defined type of sensory information, e.g., human faces. However, the bounds of “narrowly” in this context are subjective and can include a quite broad range, which is often dependent on the brain region or behavior being examined, the methodological means by which they are being examined (psychology, neuroscience, philosophy), or the particular beliefs of the researcher who is examining the topic. For example, while human faces may seem a narrowly defined kind of information, compared to all types of faces or objects, a human face would be broad relative to visual input associated with vertical or horizontal lines.

Encapsulation and domain specificity will be examined further in a later section, along with other elements of modularity that have been used to define the term. However, first a higher-level distinction of two uses of the term ‘module’ must be described, the anatomical module versus the functional module, so that I can later properly examine how modularity differs from neural reuse.

Section 2: Functional Versus Anatomical Modularity

At its most basic and fundamental level, a module is both self-contained and dissociable from other modules. However, at a more detailed look, an important demarcation is required in terms of precisely what is intended when referring to a

“module.” In one sense, a module can refer to a system that is domain specific, informationally encapsulated and performs a functionally independent cognitive process; what is termed *functional modularity* (Bergeron, 2007). A common example used to illustrate functional modules is the distinction made between Broca and Wernike’s areas, and their respective contributions to language interpretation and production. From neuropsychology, we have learned that patients with damage localized to one of these regions continue to exhibit near normal functioning of the other. Patients with damage to Broca’s area primarily are unable to produce language with less trouble understanding it, while those with damage to Wernike’s area cannot understand language but can more or less speak with proper grammar, syntax, rate and intonation. Thus, Broca’s and Wernike’s areas are functionally independent of each other.

A second framing of modularity, *anatomical modularity*, assumes that the implementation of function by most modules occurs over a localized and relatively small area of the brain (Bergeron, 2007). This understanding is largely based on neuropsychological evidence, where double dissociations are often associated with damage to localized regions in an otherwise normally functioning brain. In other words, damage to a relatively small brain region is consistently correlated with a relatively specific loss of function. Together, these two ways of thinking about modularity have framed a significant amount of thinking and experimentation about how the brain is organized and functions.

When function and anatomy are viewed in this light, a building block model emerges where modules could conceivably be assembled from other modules in a multitude of ways to support different functions. However, how does a massive

modularist researcher know at which grain size to look, or more importantly, at which level to stop digging deeper for dissociable and shareable operations? Prior to cashing this out further, it will prove useful to examine the underpinnings associated with two common perspectives on modularity: Fodorian and massive modularity.

A Comparison of Modularity Perspectives

A great divide in accounts of modularity was created when two apparently incompatible approaches to explaining how we should understand a module gained attention. *Fodorian modularity* sets out a demanding and functionally limited account (i.e., dealing strictly with a narrow range of inputs and their outputs) that does not involve central processors, whereas *massive modularity* opens the door for a more liberal and widely applicable interpretation (i.e., applying to a broad suite of inputs and higher level cognitive processes). Each perspective of modularity is useful in providing a conceptual framework for how the mind works; however both cannot be right and each has its limits. In the following sections, I will investigate these perspectives and provide a foundation on which it can be argued that neural-reuse provides the best approach to understanding the brain's cognitive architecture.

Section 3: Fodorian Modularity

In his 1983 book *The Modularity of Mind*, Jerry Fodor sets out the first detailed explanation of the mind being comprised, in part, of distinct modules that perform unique cognitive functions. Fodor sets out clear and precise criteria of a module. The list of criteria put forward by Fodor is meant to provide the rules by which one can determine if

a function is produced by a module. In general, most of these criteria are expected to be present within any given module, particularly the fourth on the list below - encapsulation.

The nine criteria described by Fodor include: domain specificity, mandatory operation, limited access, fast processing, encapsulation, shallow input, fixed neural architecture, characteristic and specific breakdown patterns and characteristic ontogenetic pace and sequencing. For the sake of simplicity, I will focus only on those criteria that are particularly useful for differentiating between fodorian modularity, massive modularity and neural reuse. These criteria include domain specificity, mandatory operation, limited access, encapsulation and shallow inputs.

I have excluded fast processing because it is implied in what is described as mandatory operation and limited access. The fixed neural architecture criteria has been excluded because it is inconsistent with what we know about the brain, in term of neural plasticity and its ability to reorganize under certain conditions, e.g., co-opting of neural anatomy by one sensory modality when not performing the function it normally performs (reading brail co-opts regions of the visual cortex), or within a sensory modality as well (phantom limbs that can be sensed through other body parts, like the face). Similarly, characteristic ontogenetic pace and sequencing has been omitted because for most descriptions of massive modularity and for neural reuse, cognitive function is largely influenced by learning and experience. I have excluded specific breakdown patterns not because it is not relevant to massive modularity and neural reuse but because it is relatively simple to explain as it refers to a module being functionally dissociable, i.e., if it is damaged, other modules or systems will not be effected.

Domain Specificity

First, a module is domain specific, which means that the types of inputs processed in a given area are limited. Hirschfeld and Gelman (1994) explain that it is “the idea that all concepts are not equal, and that the structure of knowledge is different in important ways across distinct content areas.” In other words, certain areas of the brain process specific types of information and not others. Examples of domain specificity could include voice recognition or colour detection; some have broadened the use of this term to the sensory modality level at large, however this reflects a conversation on what, precisely, “specificity” means, i.e., the relative range of a function as measured by the number of inputs that influence it. Thus, colour detection would be *more* domain specific than general vision. It’s not entirely clear where Fodor stands on this issue, however his tendency is to use highly specific examples in this context.

Bates (2001) adds some conceptual clarity by providing a list of four ways in which domain specificity can occur: 1) the task or problem to be solved; 2) the knowledge or representations that must be present somewhere in the brain of an individual who can solve the problem and produce the requisite behaviours; 3) the neural mechanisms or processors that are required to sustain those representations; and 4) the genetic substrate that makes 1-3 possible (in interaction with some environment).

Thus, domain specificity depends on cognitive responses (behaviours) to specific environmental problems (a stimuli or set of stimuli). The question of how specific a domain must be to fit this concept is one that is difficult to resolve in a modular mind context, given that what a module is and where it is located (e.g., granularity) differs from person to person. For Fodor, most cognitive processing occurs by domain-general

mechanisms, with domain specific modules limited to relatively few functions. I will argue later in this paper that this inconsistency in describing modules can be avoided by taking on a neural reuse perspective, which is less dependent on anatomical granularity and more on function (i.e., roles or operations of a collection of neurons). Over (2003) is supportive of Fodor's view of domain-general predominance and provides evidence from an evolutionary perspective to support this claim – further elaboration on this is provided in the Section 5 *Development and Evolution*.

Mandatory Operation

Second, a module's operation is mandatory for a given set of inputs. If the right type of input is provided, the operation that leads to a function will always be activated. It is quite easy to imagine this for certain functions; if the image of a dog is presented to an individual's line of sight, the individual can't actively stop the sight of this image, nor can the individual actively change aspects of the image such as its colour. Similarly, if English is spoken to an English speaker, the English speaker cannot hear the English sentence as anything other than an English sentence. Given that Fodor's perspective excludes central processing, the description of mandatory operation is straightforward. Indeed, what enables an operation to be mandatory is the absence of any intervention by a central system.

Limited Access

Third, the opportunity for central processors to retrieve information from a module is highly limited (Fodor, 1983). Indeed, Fodor (1983) has proposed that the information processed in modules is not accessible to conscious processing and thus

effectively acts autonomously and in parallel with higher cognitive functions such as critical thinking and problem solving. For example, we have no conscious access to the processes that detect faces or recognize colours. Thus, modules are not available for introspection; the representations that they produce are not open to interpretation, and they cannot be consciously altered or transformed. This criterion of Fodorian modularity clearly overlaps with that of mandatory operation, as automaticity operates by way of unconscious processes, yet can be made available to one's awareness. Thus, an individual can access the output of an operation but cannot access the actual operation in order to change the output.

While this proposition does not set off any alarm bells, it also adds little, if anything. Here Fodor uses conscious access as the bar by which he defines access, however very little of what is happening in the brain at any given time is accessible by our conscious thoughts. So what of the non-modular brain components that are not accessible to consciousness, and more importantly, what of the central processors that are not accessed by consciousness at a given time? When considered together, "modules", according to this particular criteria, share considerable ground with the brain at large, and therefore add little value to the discussion.

Information is Encapsulated

Fifth, modular information is encapsulated which means that the module cannot use information stored in other areas, although it could conceivably access information stored within its own databases (Carruthers, 2006). As alluded to in the introduction to this section, the requirement that modules are encapsulated in their processing of information is central to Fodor's definition of modularity because it enables one to carve

up the brain into functionally dissociable and separately modifiable parts; these elements are what best define modularity at a high level. As Deroy (2010) puts it “...being two ways hermetic, modules can be compared to tubes: nothing escapes before the output, and nothing gets in except at the level of inputs.”

The evidence for encapsulation largely comes in the form of illusions. A classic example comes via the Müller-Lyer illusion, for which an observer is aware that two lines are equal length but cannot prevent herself or himself from perceiving one as longer than the other. However, other illusions, such as the well know “duck-rabbit drawing”, do not produce the same effects. Indeed, with the duck-rabbit drawing, what is perceived is generally what the observer wants to perceive (Deroy, 2010). The same can be said for the famous Necker cube. For these illusions one can intentionally move from one perception of the image to another.

Deroy (2010) proposes that the difference between the Muller-Lyer illusion and the duck-rabbit/Necker cube illusions does not challenge Fodor’s encapsulation requirement as the effects of the duck-rabbit and Necker illusions are a result of the order in which particular components of the images are attended to, i.e., the input changes depending on where one’s attention is drawn first, rather than the information being not encapsulated. In other words, the ability to switch back and forth in the duck-rabbit and Necker cube examples occurs at the pre- or post-module level; context is important (Deroy, 2010).

It is important to understand the distinction between encapsulation and *Limited Access* because, at times, they can easily be confused for the same thing, when in fact they are very different. Whereas *Limited Access* means that higher brain areas cannot

access processed information within a module, an encapsulated module (or information contained in modules) cannot access information from other brain regions for further processing. Encapsulation is the most fundamental element of Fodor's description of modularity, and for this reason it has also been the most studied and criticized element. In the next section I will provide some of the most compelling arguments against encapsulation, most of which have been developed by leading massive modularists.

Shallow Inputs

Sixth, modules have what Fodor refers to as "shallow inputs", which means that information processed by the module is limited in the kinds of information it processes and the amount. As Deroy (2010) describes Fodor's definition of shallow inputs, "representations generated by modular, and noticeably perceptual systems, are somewhat primitive, pre-semantic, non-interpreted and un-appraised." Another useful descriptor used by Deroy (2010) for 'shallow' is the idea that this output is pictorial and non-conceptual. Carruthers (2006) also highlights that Fodor's inputs to modules are shallow in that the information they contain does not generate thoughts or beliefs in the module, e.g., visual input to a model might produce surfaces and edges but not the recognition of a particular object. Above all, shallow inputs are uncomplicated and taken in quickly from the environment.

In further dissecting the above characteristics, Deroy (2010) raises some difficulties in determining where to draw the line. For example, Deroy (2010) questions whether the level at which colour and shape information are processed together, or as Deroy puts it, "when colours and shapes are bound together and attributed to the same location." Deroy (2010) suggests that this level of investigation might not be shallow

enough to be modular processing; rather that it is a product of two modular outputs. Consequently, Deroy (2010) challenges philosophers to develop better terminology to more accurately reflect what type of shallow representations form from low-level, automatic, encapsulated perceptual processing. Indeed, this is another example of difficulty in identifying the appropriate grain size at which to define a module. A full examination of granularity can be found in Section 11.

Section 4: Massive Modularity

Developing a framework for massive modularity is not as straightforward as for Fodorian modularity; this is not a surprise given the specificity and narrow constraints required under the latter. But how flexible is massive modularity, relative to Fodor's view? It is conceivable that one answer to this question could employ an extremely simplistic definition of what constitutes a module, such that the brain can be broken apart into hundreds or even thousands of components that each have their unique function and any of whose elimination or alteration would not prevent the general functioning of the whole, i.e., they are functionally dissociable and separately modifiable. This would resemble an extended version of Fodor's modularity; one that explains a much broader number of functions using the same general principles. But what happens when we dig deeper? For example, how does massive modularity deal with encapsulation and domain specificity, the two fundamental criteria of Fodor's module?

Where Fodorian modularity is confined to the periphery in terms of the inputs it considers and the functions it supports, massive modularity is reaching farther in. Indeed, the massive modularist brain is comprised of a suite of solutions that have evolved to respond to environmental problems presented to our ancestors and the lineage of common

ancestors before them. While massive modularists generally accept this basic principle, I will explain in this section how and why the massive modularist community has become fragmented by a number of variables relating to the properties of a module. The first of these differences that I'll discuss involve the types of function that should be considered outputs of modules. The second will look at whether modules can share their component parts, i.e., one or more of their operations.

Modules are often referred to as Darwinian modules as they are supposed to have been formed by the process of natural selection (Machery, 2007). So, in the way that organs and elements of the body have evolved to enable organisms to adapt and survive in their environments, so too have distinct modules within the brain. While the brain is thought to mainly consist of these Darwinian modules, humans also have non-modular systems that respond generally to perform a variety of tasks (Machery, 2007). Machery (2007) highlights reading ability as one such non-modular system, noting that reading is a relatively recent cultural invention that makes use of a collection of modules that evolved for other reasons.

The distinction being made by Machery between Darwinian modules and non-modular systems is meant to differentiate between a function that is carried out by an evolved module and by a non-modular system that has not evolved but rather has been established by learning or development with what already exists within the brain. Ultimately, this distinction is tied to the assumption that a module cannot be comprised of sub-modules. This leads us to an important question – if a collection of modules work together to produce a function, can this be considered a non-Darwinian module? In other words, can the reading function be understood as a module, yet one that is non-

Darwinian? The answer to this question has created a divide within the massive modularity community.

Sperber (1994) suggests that modules can be combined to create a higher-level module. Sperber (1994) describes hypothetical creatures he calls ‘protorgs’. Protorgs have cognitive capacities that enable them to hear noises and feel vibrations made by larger animals, each processed by a unique module (Sperber, 1994). In order to maximize the efficiency of these modules, it is important for the protorg to initiate a flight response only when both modules are activated simultaneously, because a larger animal is present only when these two inputs occur together (Sperber, 1994). Sperber (1994) refers to the mechanism whereby both inputs are integrated as a module; therefore the noise recognition and vibration recognition modules together form a “problem solver” module. Sperber (1994) also proposes that this module is domain specific, which Blackmore (2013) highlights as an example of functional specificity being mistaken for domain specificity. Indeed, Sperber’s module cannot be domain specific if it is activated by two different types of input (noise and physical vibrations). So, we now have Machery’s module, which acts alone to support an evolved function, and Sperber’s module, which can be comprised of smaller modules but need not be. This is not the end of the ambiguity concerning massive modularity.

Peter Carruthers, a philosopher of mind who has been influential in articulating proposed characteristics of massive modularity, describes modules as being,

Functionally dissociable, intentionally characterized processing systems, each with its own neural realization. Modules need not be encapsulated, domain specific or innate (although many probably are). And the neural systems that realized them certainly need not be anatomically localized. On the contrary, modules can be realized in spatially distributed interconnected networks of brain

regions. Moreover, many modules are constructed out of, and share parts with, other modules (Ritchie & Carruthers, 2010, p.289).

What is most interesting in the quote above is that Ritchie & Carruthers permit not only a combining of modules to support functions, but the sharing of a module's components with other modules. This view of massive modularity is inconsistent with the views of Machery (2004), for example, as it includes non-Darwinian modules, and also opens the door to domain generality through the sharing of sub-components. An analogy might be useful in describing how Carruthers's view differs from other massive modularists. Consider a dinner plate that consists of macaroni and cheese, salad, and apple pie, where each element on the plate is a module that supports the function of a meal. Massive modularists like Machery would state that each of these three foods can be recombined with other foods to create different meals. So, in this view of modularity, there are dissociable parts (three modules) that can be reorganized to support different functions. Thus, the module is the lowest level of decomposition, i.e., there are no sub-modules. Likewise, Sperber would agree that each of these modules can combine with other modules, but in contrast with Machery, he would describe the combination of each group of modules (i.e., each meal) as a module.

Now, what Ritchie & Carruthers has permitted is that each food (module) in the meal can be comprised of shareable components. So, for the macaroni and cheese, the macaroni can be combined with a component from another module (say, the meat sauce from a lasagna) to form a new sub-food (macaroni and meat sauce). Likewise, the cucumbers in the salad can be combined with a lunchmeat and bread to form a sandwich. To clarify this analogy, Sperber's module can combine with other modules to produce a higher-level module but it seems that whatever neural circuitry lies within each module

(sub-module) supports only that module (or sub-module). According to Carruthers, each module can combine with other modules, and also share its parts to be used in other modules. So, for Carruthers, modules can be built out of parts of other modules.

Carruthers has introduced the term ‘wide scope’ encapsulation to describe his view of modularity, which differs from Fodor’s narrow scope encapsulation. Whereas Fodor’s encapsulation limits processing to that of domain-specific stimuli and within-system databases and excludes any other information, Carruthers proposes that an encapsulated system “*can’t* be affected by *most* of the information held in the mind in the course of its processing” (Carruthers, 2006). We can also take this to mean that although some components of modules can be shared with other modules, this is not common. Carruthers (2006) underscores that wide scope encapsulation is frugal, just not as frugal as Fodorian (narrow) encapsulation, where nothing gets in; it uses heuristics and stopping rules to determine how much information is accepted. The expanded scope Carruthers invokes for encapsulation is necessary if massive modularists are to hang on to this concept because we know that central processes are not encapsulated in the way that Fodor has described.

Carruthers (2006) refers to the mind-reading module as one that involves wide-scope encapsulation. Here it is thought that the mind-reading module must query other systems in order to effectively predict what another is thinking or what his or her next behavior will be. Sperber (2001) uses the example of a general modus ponens module (if P then Q) as a central processing module, as it is intuitively a strong candidate for a cognitive reflex, as it does not require querying other databases for its output. However, Fodor has claimed that such a modus ponens function, particularly one that is generally

applied to reasoning across many topics (e.g., numbers, food, people, plants, etc), would not be encapsulated or domain specific (i.e., numbers and people would be distinct inputs) and therefore could not be a module (Sperber, 2001). Here is one example of how Fodor's strong interpretation of encapsulation can create problems for massive modularists.

As alluded to in the previous paragraphs, the most significant divide between the two approaches to modularity can be seen in the types of output that are produced. While Fodorian modularity permits only rudimentary outputs, such as face recognition, massive modularity includes an account of a vast number of cognitive functions that are processed centrally (Carruthers, 2006). Extending this architecture to include such high level processing necessarily contradicts certain elements of Fodor's perspective. First, it is conceivable that some modules are turned on by multiple beliefs, such as a module for logic (Carruthers, 2006); this is clearly incompatible with Fodor's domain specificity. However, it is conceivable that some modules are in fact domain specific.

Second, at face value, Fodor's fast processing criteria also appears to be incompatible with massive modularity, as 'fast' here is relative to more complex and temporally demanding central processing, which is an acceptable level of processing, which at first thought would seem inconsistent with the processing of higher-level functions (Carruthers, 2006). However, it should be noted that there remain doubts that central processing is in fact slow. For example, Gigerenzer & Todd (1999) propose that some reasoning and decision making is carried out rapidly through inferences acquired using fast and frugal heuristics.

Third, it is not clear whether massive modularity is consistent with Fodor's criterion of innateness as the massive modularist community has not yet settled on a position on this topic. Indeed, this failure to reach common ground on many elements within the massive modularity community presents a significant problem with the perspective generally, and will be discussed in more detail in Section 6). While Fodor would likely assert that both evolution and development play a role in module formation, Tooby and Cosmides (1992) contend that through natural selection the brain has been molded into a collective of modules in response to past evolutionary pressures. For Tooby and Cosmides, modules appear to be fully hard wired. Conversely, Paterson et al (1999) have proposed that modules are formed through learning and development. Development from birth to one year old is associated with myelination of neurons, beginning with neurons in the base of the brain at birth, and terminating with myelination of the occipital, parietal and frontal lobes at 8-12 months (Paterson et al, 1999). It has been suggested that the rate at which myelination occurs during this timeframe can effect acquisition of cognitive abilities, i.e., increased myelination yields increased cognitive capacity (Paterson et al, 1999).

Based on Carruthers's interpretation of massive modularity, it appears that many of Fodor's criteria would need to be weakened or completely removed in order to align with a massive modularity perspective. Carruthers (2006) summarizes the notable differences and compatibilities between the two perspectives in the following quote:

...the properties of having proprietary transducers, shallow outputs, fast processing, significant innateness or innate channeling, and encapsulation will very likely have to be struck out. That leaves us with the idea that modules might be isolable function-specific processing systems, all or almost all of which are domain specific (in the content sense), whose operations aren't subject to the will,

which are associated with specific neural structures (albeit sometimes spatially dispersed ones), and whose internal operations may be inaccessible to the remainder of cognition (p. 12).

Clearly, what massive modularists call a ‘module’ is not what Fodor means by this term. Furthermore, what some massive modularists call a module shows almost no resemblance to what other massive modularists call a module (e.g., Machery, 2007; Sperber, 1994, & Carruthers, 2010) – this will be discussed in detail in Section 6. I have focused on Carruthers’s description of modularity in this section as it is the most relevant to identifying the problems with Fodorian modularity and it is also the most useful in examining the merits of neural reuse by contrast. In the following section I will examine the evidence from evolution and development to compare the merits of multiple interpretations of modularity and modularity itself. I will then be positioned to critique modularity in Section 6.

Section 5: Evolution, Development and Learning: Support and Pitfalls

For any examination into the functioning of the brain, whether it be from a Fodorian modularity, massive modularity or neural reuse perspective, identifying evolutionary pressures or developmental mechanisms that could have influenced cognitive architecture will do two important things: 1) it will increase the explanatory value of a perspective by providing a conceivable rationale for why the brain functions in a certain way; and 2) it will provide a framework under which multiple perspectives can be assessed against one another. Ultimately, any theory of mind must be compatible with evolutionary theory. The section below describes the evolutionary and developmental evidence for massive modularity. In a similar vein, the rationale for neural reuse under the same framework will be examined in Section 9. Fodorian modularity will not be

examined in this context, as this would not add significant value to what is being considered with respect to the other two perspectives.

The centre of the massive modularity argument lies in its ability to explain how modules form, in both developmental time and evolutionary time. In fact, this perspective has the most at stake on this particular topic, as evolved responses to historic environmental challenges are what have driven the massive modularist perspective from the start.

The increased range associated with a massive modularity perspective (i.e., it's application to functions that are accessible to consciousness) invites a discussion of the relative importance of genetics versus learning for a given function. Over (2003) emphasizes that massive modularity should not be reduced to genetic determinism. However, substantial evidence suggests that genes have played a leading role in it.

Evolution

Carruthers (2006) describes a module as a functioning system that has been assembled from sub-components that each play a role in the module's overall function. This characterization is well established in biology, across many levels (e.g., gene, molecular, organ, organism) and has been co-opted by massive modularists to explain the composition of modules in a similar vein. A frequently cited example of an evolved function is fear of snakes (Confer et al, 2010). It is not hard to imagine that being fearful in the presence of snakes provides an evolutionary advantage as it results in avoidance behavior and therefore reduces the probability of dying from a snakebite. Fear of snakes is an adaptation; it occurs automatically, is hard to turn off and it is associated with specific neural circuitry (Confer et al, 2010). However, this function is also responsible

for the by-product of producing fear when confronted with non-venomous snakes as well (Confer et al, 2010). This raises an obvious question; how can we determine the level of granularity of a module, e.g., is there a module for each fear (snake, heights, loud noises) or simply one for fear in general, and how is this granularity determined?

Evolutionary psychologists are in the business of investigating and hypothesizing selective pressures that may have led to the development of a particular cognitive function. The cognitive module is a useful tool for exploring function by natural selection. However, it is conceivable that evolutionary psychologists are biased in identifying their epistemological starting point. For example, a massive modularist who sets out to determine the evolutionary origins of face recognition has locked herself or himself into the assumption that a module exists for face recognition, thereby limiting the evidence that they are willing to consider and predisposing them to overlook other relevant information. Anderson (2014) sums up this problem as follows:

One thing that is important to note is that the evidence for functional diversity in the brain becomes clear only when one looks across multiple experiments to characterize local function. For a modularist, this is not a natural scientific strategy to adopt, for if brain regions are dedicated to individual task domains or stimulus types, then once one has identified the stimulus or task that effectively causes a regional response, one need look no further to be in a position to characterize function (p.16).

Likewise, Bechtel (2003) claims that many evolutionary psychologists assume the brain is modular at a coarse grain, when in fact they should be exploring cognitive function at a much finer grain. Bechtel points to an economist and a computer scientist, who have described the functioning of computers as modular, to help cash out this claim: Herbert Simon and David Marr. He notes that Simon realized that the modular approach computers used to solve problems (i.e., breaking complex tasks into sub tasks) could

apply to evolved human systems, such as the brain (Bechtel, 2003). In fact, Simon was of the view that evolution could only have produced complex systems that are nearly decomposable (Bechtel, 2003); a system comprised of modules that don't share parts. Similarly, Marr (a physiologist turned computer scientist) developed models for physiological systems by employing what he called a *principle of modular design* (Bechtel, 2003). The principle of modular design describes components within a system as being weak and having little effect on each other, meaning that individual components can be identified and examined independently (Bechtel, 2003).

Bechtel (2003) acknowledges that the ideas of Simon and Marr are logical to a first approximation, however he underscores that in many systems the component parts do not operate independently of each other; they often interact. Moreover, the operations of the component parts often cannot be easily understood vis-à-vis the function of the system. For example, the component parts of the fermentation process do not carry out sub-level fermentations, rather oxidations, reductions, phosphorylations, etc. (Bechtel, 2003).

Bechtel (2003) argues, rightly I would suggest, that neuropsychologists do not typically examine the function of a particular neuro-anatomy at the appropriate grain size; information-processing operations should be examined at a much finer grain. However, this is an understandable and expected corollary of the technology by which we view cognitive function in action. The sophistication of neuroimaging technology, such as PET and MRI, is rudimentary relative to the grain size that could conceivably be imaged. This has required researchers to explore macro-level functions that are not particularly useful in shedding light on the functions or sub-functions that were actually molded by

selective pressures over time. Thus, in their defense, neuropsychologists have been limited in the territory that they cover.

Bechtel (2003) also proposes that operations can be recruited in a number of ways to support multiple functions. It is this interpretation of modules by Bechtel that may have set the groundwork for neural reuse. Bechtel highlights the importance of connectivity patterns within the brain as an effective tool for gaining insights into its fine-grained operations. A section on connectivity can be found in Section 16, particularly in relation to the fusiform face area. Taken together, Bechtel's work demonstrates that if one is using an evolutionary framework to assess the validity of a given approach for understanding cognitive architecture, it is important to consider which functions have evolved, rather than been co-opted, and the grain size of neural circuitry that is contributing to a function of interest.

In a similar vein, Confer et al (2010) suggest that adaptive behaviors are associated with modules, however not in the Fodorian sense. Rather modules "share components and interact with each other to produce adaptive behavior", and therefore cannot be encapsulated. Confer et al (2010) underscore the importance of examining both ultimate and proximate explanations when considering a particular function of a psychological mechanism - the former being the evolutionary cause and the latter the workings of the mechanism itself. Indeed, having the explanation for one often provides hints about the other. For example, our understanding of stranger anxiety in infants would be incomplete if we didn't know that the function is to avoid potentially dangerous humans (ultimate explanation) or that this function is associated with increased crying,

which develops at 6 to 8 months and is more intense around males (proximate explanation) (Confer et al, 2010).

The premise upon which development is thought to be relevant for modules is that the differentiation of modules is governed by developmental/genetic switches, each causing the formation of components of a given phenotype (Carruthers, 2006). A set of genes is responsible for any differentiation and therefore susceptible to selective evolutionary forces (Carruthers, 2006). Thus, according to this view, modularity is “an extremely common solution to evolvability” (Carruthers, 2006). Essentially, evolution can act on components within modules or can act at the module level in relation to other modules. Together, it is not difficult to imagine the vast opportunity for nature to select favourable properties of the human mind to increase fitness over time. However, while modules can be thought of as distinct for cognitive purposes, they do not act in isolation; behaviours of an organism over time, across generations, and ultimately the fitness of an individual or lineage, is defined by a collective of cognitive events. For example, better vision might be selected for only in conjunction with a level of motor control that can enable the organism to take advantage of better vision, whether it be hunting, evading a predator, etc. This must always be kept in mind when contemplating evolutionary effects.

While understanding the evolution and learning behind a function can be insightful and provide explanations of why and how we respond to stimuli the way we do, it is conceivable that a number of brain functions are spandrels – a brain region that has evolved for adaptive reasons for one function but has been exapted to perform a new function. Carruthers (2006) proposes that language is an example of this type of

functional occurrence. Given that many of the neural systems required for language are common across several species, it is possible that our ability for language is the result of a very small change, such as a single random mutation, or an exaptation of a change that was selected for another purpose (Carruthers, 2006). Terrence Deacon (1997) expands on Carruthers's general claim that language is not attributed to an evolved module but has resulted from the brain making new use of neural circuitry that has evolved for other reasons by hypothesizing an operation that this co-opted circuitry might perform.

In his book *The Symbolic Species* (1997), Deacon makes use of neurobiology, evolutionary theory, linguistics and semiotics to demonstrate that language could have co-evolved with the brain. Of particular interest is Deacon's suggestion that there exist significant problems with what we think we know about the origins of language. Deacon (1997) requests that we set aside our thinking that language is largely nature driven (mentalese), largely nurture driven (simple associationism) or a partnership between the two, as this issue is secondary. What matters more, Deacon proposes, is that human language is an anomaly, not comparable to communication of other species, just as facial expressions of humans are not comparable to language (Deacon, 1997). However, Deacon (1997) further clarifies that the natural corollary to this – the idea that humans possess a unique, language-producing neural circuitry (e.g., Chomsky's universal language organ), is erroneous. Moreover, the acquisition of language by children lacks a physical mechanism – an explanation (Deacon, 1997). In other words, “universal grammar” is no more than a placeholder for what is yet to be learned (Deacon, 1997).

While it is certainly conceivable that a specialized area of brain matter may have evolved and is well positioned for language, understanding the underlying operations that

are carried out by components of this area may be what's needed to add clarity to what it *means* to be a "language area." Similarly, this would be equally relevant for the language by way of greater intelligence argument; a more broadly distributed (anatomically) unspecialized approach to language acquisition. Deacon (1997) proposes that it's not the supposed complexity of language that presents a limiting factor for other animals; he explains that language is not complex as simpler forms (e.g., the speech of a young child) have not evolved in other animals, and other animals, such as birds, have comparable forms of communication, in terms of their complexity, to that of children. Deacon (1997) has proposed that function of *symbolic reference* is more likely to provide a roadblock for the development of language in other animals. Deacon (1997) explains that an explanation for the evolution of language must be rooted in understanding *symbolic reference*. In other words, symbolic reference could support functions beyond that of language. Moreover, developing a better understanding of the component operations that drive symbolic reference is sure to shed light on its origins.

An experiment conducted by Clune et al (2012) that employed computation simulations of differing evolutionary contexts support the claim that modularity could have evolved as a means to reduce connection costs within the brain, which can be very energy-intensive. In particular, Clune et al (2012) revealed that both modular and non-modular architectures could be high performing, however modular systems demonstrated greater evolvability as this architecture enables sub-components to be altered without disrupting other systems. As expected, the results confirmed that modular architectures adapt better to changing environments (Clune et al, 2012). This further supports Marr's principle of modular design, which was discussed earlier in this section. However, these

same efficiencies could be gained from a brain comprised of specialized workings that are domain general, i.e., neural reuse (this will be discussed in more detail in Section 9). A particularly important element here is the level of granularity at which one defines a module. This topic will be expanded on in Section 11, vis-à-vis Bechtel's "decomposition of the performance."

Development and Learning

Now that modularity has been examined in an evolutionary context, it can be considered against what we know about development. While the primary objective of this section is to determine how well modularity fits within the rules of evolution and development, I am also attempting to gauge the influence of each effect on the formation of modules, i.e., to what extent evolution plays in the formation of modules, compared to that of development. Based on what we know about the mind, it is unlikely that the brain's cognitive architecture is formed by only one of these processes. But if both are involved, what is the relative influence of each on cognitive architecture? Answers to this question can not only help us to better understand what is meant by massive modularity, but also how massive modularity compares to other perspectives on cognitive architecture.

The evolution of a species and the genetic composition of individuals provide the hardware for a limited range of potential behavioural responses to a given environmental input. Confer et al (2010) provide examples of proposed explanations for learning mechanisms that are thought to be linked to evolved modular neural machinery, rather than a general learning adaptation. Learned incest avoidance is clearly important for avoiding undesirable reproductive partners. There is evidence that suggests that co-

residence with a member of the opposite sex during development acts as a cue for genetic relatedness and leads to incest avoidance activation (Lieberman et al, 2003). Another example is food aversions, which seem to use completely different learning mechanisms than those invoked for incest aversion (Confer et al, 2010). Confer et al (2010) explain that humans have evolved mechanisms to first recognize correlations of ingested foods and nausea, and second, to recall this correlation in the future when making decisions about food intake. Finally, Confer et al (2010) refer to a third learning mechanism for recognizing status and prestige of others. This mechanism acts by causing individuals to recognize which individuals within the group receive the most attention and to replicate their qualities (i.e., prestige criteria) (Confer et al, 2010).

Following on this, Paterson et al. (1999) found that the development of various cognitive domains, including rapid auditory processing, face processing, object permanence and joint attention, is dependent on the recruitment of new brain regions as the individual matures and their performance is strengthened and expanded (Paterson et al, 1999). For example, early in development, face processing is dependent on subcortical regions but as the child develops, other areas, such as the fusiform area, are recruited to provide a more sophisticated function (Paterson et al, 1999). However, brain regions associated with the acquisition of a function in development are often not the same as those associated with maintenance of the function in adulthood (Paterson et al, 1999). For example, joint attention is at first achieved in children via the visual cortex but as they develop, activity associated with this function is largely seen in the frontal cortex as the child acquires the ability to engage others (Paterson et al, 1999). This is important, as it demonstrates that the brain is plastic and that at least some of the

specialization it acquires is largely influenced by the type and quantity of input it receives. This evidence aligns well with the neural reuse perspective because it provides a point from which we can imagine certain regions developing to support many different uses based on a common strategy, e.g., the hierarchical sequencer thought to be the working of Broca's area.

Much as with face processing, joint attention becomes more refined with development, leading to a more effective cognitive function. The examination of function from a developmental perspective provides a good basis from which to imagine a compiling of sub-functions to give rise to higher-level functions. While the above findings are largely consistent with a massively modular description of cognitive function, they don't require one. Once this picture is drawn, it becomes easy to imagine a cognitive system in which sub-level workings are reused and reassembled to produce a suite of cognitive functions. I would suggest that the idea that a module for joint attention moves from the visual cortex to the frontal cortex over development is less plausible than premature neural workings recruiting and combining with other neural workings, over time, to establish a more mature joint attention-performing ensemble, whose function might also be employed (reused) for other uses. I say this simply because the movement of a module from one location of the brain to another would seem inefficient as resources would be required to establish an anatomical region that will ultimately be dismantled. In addition, while it is easy to imagine modules as anatomical units that establish themselves only after development, it seems unlikely that they would be encapsulated or would be domain specific under these conditions. It would seem to contradict the plasticity we see throughout development, and beyond.

Karmiloff-Smith (1992) does respond succinctly to this particular challenge. She proposes “...the mind becomes modularized as development proceeds.” In other words, the heightened plasticity of the developing brain responds to environmental cues and “with time, brain circuits are progressively selected for different domain-specific computations.” Karmiloff-Smith (1992) suggests that the brain is predisposed to form certain modules but in an epigenetic sense, rather than the hard nativism Fodor has proposed. More specifically, “...nature specifies initial biases or predispositions that channel attention to relevant environmental inputs, which in turn affect subsequent brain development” (Karmiloff-Smith, 1992). We see here what we have seen throughout our examination of modularity; an attempt to further make modules, or semi-modules, less constrained and, in fact, less modular.

Karmiloff-Smith (1992) has emphasized that learning new skills at first requires significant conscious thought and attention but with practice these skills become automatic. However, the perfected skill is not to be confused with a Fodorian module as the skill is based on learning through experience and is not fixed anatomically nor is its information encapsulated. While learning is compatible with the more flexible exchange of information associated with the massive modular view, I will also show in the following chapter that learning and perfecting cognitive skills is better aligned with a neural reuse perspective.

There are compelling evolutionary and developmental arguments favouring a massive modularity perspective of neural architecture, e.g., energy efficiency associated with adapting or developing new functions, although many of these arguments also support specialization which does not require a modular architecture (neural reuse is also

a proposed architecture centered on specialization). Indeed, the same evidence supports a neural reuse perspective on cognitive architecture. In fact, I will describe in Section 9 how other evidence provides greater support for the idea that evolution has produced a brain that is not comprised of modules, but one that has evolved (specialized) workings that are shared and reused to support our repertoire of cognitive functions. However, I will first identify some problems with Fodorian and massive modularity perspectives of cognitive architecture.

Section 6: A Critique of Modularity

While the previous section identifies a number of advantages for massive modularity in an evolutionary and development context, there also exist a number of challenges faced by modularity, in general, as it relates to a cognitive architecture of the brain. As Barrett and Kurzban (2006) point out, the strengths of Fodorian modularity are also its vulnerabilities. The narrowness of Fodorian modularity enables it to be particularly useful as a psychological concept, however given that Fodor assumed that most cognitive functions could not be explained, his perspective is very limited in terms of its potential application. In addition, as research in cognitive science accumulates we are beginning to understand the central process that Fodor claims cannot be understood.

A popular argument against Fodorian modularity is that it is often incompatible with our current understanding of cognitive architecture, i.e., processing within a module is influenced by external processes and not simply entrance inputs (Barrett & Kurzban, 2006). This is targeted at the crux of Fodor's view; that modules are encapsulated. And it is likely the reason why Fodorian modularity has morphed into the less confined massive modularity.

Prinz (2006) identifies a number of problems with encapsulation. First, he attempts to demonstrate that conclusions drawn by many researchers regarding our perception of visual illusions are mischaracterized and, in fact, do not provide good evidence for encapsulated modules. In particular, he explains that our inability to view the Muller-Lyer illusions differently after we know that the lines are the same length could be a result of a cognitive heuristic wherein a perception always trumps a belief when the two are in conflict (Prinz, 2006). Prinz (2006) refers to the Necker cube and the duck rabbit as examples of illusions that do not produce a conflict and can be manipulated by beliefs, i.e., top down processing.

In addition to top-down evidence against encapsulation, Prinz (2006) underscores that there is considerable evidence to suggest that input systems can communicate with each other, such as V.S. Ramachandran's technique for relieving phantom pain in amputees: this technique employs a mirror reflection to position the image of a limb at the site of amputation, enabling the individual to scratch or sooth the limb with positive results. It would seem here that cognitive areas that are normally influenced by sensory input are being altered by visual input.

However, these criticisms focus almost exclusively on hard Fodorian elements of modularity such as automaticity and encapsulation, and therefore don't hold as much, if any, weight in weakening the concept of massive modularity (Barrett & Kurzban, 2006). Indeed, the willingness to transform the impenetrable information pipelines of Fodorian modules to something more porous, as espoused by supporters of massive modularity, makes problems less obvious and requires a different type of challenge.

Arguments have been targeted more specifically at massive modularity. Perhaps the most common of these is the existence of domain-general abilities such as logical reasoning, working memory, etc. In particular, Fodor (2000) insists that some modules described by massive modularists accept too many inputs to be considered modules, such as a mechanism that distinguishes between shapes (e.g., triangles and squares). However, Barrett (2005) responds to this challenge by noting that Fodor's 'pipeline' concept of information processing does not hold in all cases, as Fodor insisted. Barrett (2005) proposes that certain processors can access a broad range of inputs yet process only a subset of them; those that fit with the particular processing system. Barrett & Kurzban (2006) note that with respect to working memory, our account becomes more detailed when one considers the subsystems that have been proposed to explain it, such as the visuospatial sketchpad, phonological loop and episodic buffer. All of which have very specific representational formats.

It seems as though Barrett and Kurzban are stretching the reach of domain specificity, in terms of how it can be reasonably defined, which is ultimately what the domain-general argument responds to. Clearly, moving to a finer grain of neural circuitry will at some point generate some sort of domain specificity, but is defining domain specificity in this way really useful? If massive modularists are content with digging deep until specificity is unearthed, it might be time to cut the line and let this particular element sink. Prinz (2006) puts it well by highlighting that if unused cells of the "visual system" can be used for touch, than even for this basic sensory modality domain specificity seems too strong. Similarly, Prinz (2006) notes that Broca's area – a brain region in the temporal lobe associated with language production – contains mirror

neurons that play a role in the recognition of manual actions and are located in a number of other regions across the cortex. It is unclear how massive modularists would deal with micro-level components, such as mirror neurons, that provide the same function across modules. First, it would suggest that the same micro-function evolved in many different areas independently, and second, it would be difficult to imagine how we could reconcile domain specificity in this case – would each group of mirror neurons respond to their own unique domain? Or would each domain activate multiple groups of independently functioning mirror neurons?

Barrett and Kurzban (2006) discuss how our ability to do cognitive work with novel stimuli has been presented as a supposed problem for massive modularity. This problem is put forward as follows: if modules evolved according to contexts experienced by our ancient ancestors, the same modules should not be equipped to process novel information, i.e., information associated with the present day that could not have existed in the past (Barrett & Kurzban, 2006). However, Barrett and Kurzban (2006) claim that novel stimuli are capable of recruiting modules that evolved under different conditions to perform different functions, because they can take advantage of the operation involved. For example, “collision-avoidance systems could be recruited in driving, strategic social cognition systems could be recruited in chess, and systems evolved for identifying objects such as tools or animals could be recruited to identify letters or words in reading”. This defense comes dangerously close to sliding over into the framework of neural reuse, as all one needs to do is move down a level and suppose that, according to the same principle, neural components of a ‘module’ can be recruited for multiple uses. This would present a problem for the module in that its parts are flexible, and furthermore, it

suggests that the module itself could be flexible as well, i.e., supporting multiple uses. It would be hard to imagine these conditions existing without contravening the rules of domain specificity and encapsulation.

Perhaps the most troubling problem with massive modularity is the inconsistency with which its practitioners describe it. The Machery, Sperber and Carruthers descriptions of a module are incompatible with one another, and each carries its own unique baggage. For Machery's module, the implication of modules not being permitted to combine with other modules (i.e., Darwinian modules are the only modules) is not considered, thus we either have a dilemma of grain size, or the brain is comprised of few modules. At one end of the spectrum, modules can consist of relatively large brain areas, however this would limit the number of modules that the brain can support as its mass is finite. At the other end, we can have very small modules that each perform a unique function but this would appear to be inconsistent with the high level of connectivity that has been observed between brain regions.

A second problem with the Machery module is we are left wondering which functions a (Darwinian) module performs and which are performed by non-modular systems, with no reliable criterion that can be applied on a case-by-case basis. For Machery, all modules are Darwinian and thus are hard wired. Machery (2007) uses the example of reading ability as being carried out by a non-modular system because we can date the inception of this function back to recent times; too recent to be the result of an evolved module. While this logic is good, for most functions it would be difficult to determine if they arise from modular or non-modular architectures, and there are not many functions to which this can actually be applied. If Machery does contend that most

of the brain is non-modular, at least he would avoid the problem of grain size discussed above, but it is not evident that this is the case. Another obvious question is why Machery excludes the possibility of non-Darwinian modules.

Sperber (1994) avoids the limited brain matter problem by allowing modules to combine with other modules to form new modules. This provides more flexibility than that of Machery's modularity as it is not constrained by brain mass, i.e., the combinatorial potential of a small number of modules can support a very large number of functions. However, it appears that once Sperber's module has been defined, all brain matter within the module supports only that module, i.e., there is no sharing of parts. Therefore, a module can contain dissociable modules within it but none of the neural circuitry can support other high-level modules.

For example, if the temporal parietal juncture (TPJ) were proposed to be a theory of mind module, neural circuits that perform operations within this region would not be available to support other functions. Thus, if the theory of mind capability of an individual were to somehow be turned off, this area would be void of any meaningful activity. This seems unlikely given that the TPJ is active in many tasks that are not associated with theory of mind. In addition, the TPJ is comprised of a relatively large brain region. Surely, if this size were attributed to modules in general, there would be room for only a small number of functions.

As discussed above, Sperber's problem solver module contains two modules, a voice recognition module and a vibration recognition module, which can act independently or together to support the problem solver module, but cannot support another higher-level module, i.e., Sperber's module cannot share its parts with other

modules. It is conceivable that the voice recognition module (or vibration module) is comprised of lower level modules, but again, these lower-level modules would act independently or support voice recognition, i.e., these lower level modules would not support the vibration recognition module in any way. For Sperber's problem solver module, it seems that domain specificity must be given up, as the voice recognition and vibration recognition involve two different types of inputs. If we were to apply Deroy's tube analogy to Sperber's module, it would appear that encapsulation also must be surrendered along with domain specificity. If two tubes (one for each of voice recognition and vibration recognition) are merging into one (problem solver module), the problem solver module is accessing information from multiple inputs.

Carruthers's (2006) view that parts of a module can be shared with other modules avoids the grain problem and non-Darwinian modules problems faced by Machery and Sperber but he commits himself to an even greater problem, distancing himself from any semblance of a "module". While Sperber must give up domain specificity and encapsulation within modules, it seems as though he has attempted to maintain a weakened interpretation of these concepts between higher-level modules, e.g., between modules that don't combine with other modules to form higher-level modules. While Carruthers's wide scope encapsulation might, at one time, have been considered similar to that described by Sperber, his remarks in Ritchie & Carruthers (2010) deviate from this watered down view and go as far as to state that modules need not be encapsulated at all. Likewise, Ritchie & Carruthers (2010) accept that modules can be domain general. So, we are left with a scenario where modules can be domain general and share information

with other modules. This appears to be a significant change in position from Carruthers (2006).

I hope I have sufficiently described the diversity of perspectives when it comes to modularity. Indeed, the discrepancies between Fodorian and massive modularity are equally matched by those observed within massive modularity. These discrepancies are not a bad thing; they represent progress. I would argue that neural reuse provides the next step in this progression; a better explanation of cognitive function than that of the two forms of modularity described above. However, as alluded to above, Carruthers' view that modules need not be encapsulated nor domain specific yet are functionally dissociable and separately modifiable requires special consideration in the context of neural reuse and therefore will be examined in greater detail in Section 12.

Chapter 2 – Neural Reuse

Modularity, both as defined by Fodor and by the massive modularists, is a useful starting point for conceptualizing the brain, both anatomically and functionally. Both perspectives address the obvious need to contemplate specialization of function when proposing a global cognitive architecture. However, it is important to remember that the brain's cognitive architecture can be specialized without being modular. In other words, the brain can pose specialized workings that are neither encapsulated nor domain specific. This is precisely what is being proposed by neural reuse perspective.

The constraints that modular perspectives place on our examination of the brain are significant and distract from what is most important. Research from psychology, philosophy, cognitive science, neuroscience and many other disciplines has contributed new insights into how the mind is organized. Indeed, advances in neuropsychology have

linked certain functions with relatively well-defined brain regions through double dissociations, such as those thought to be responsible for speech production and comprehension. Further to this, the neural circuitry for certain functions, such as vision, has been mapped out in great detail. Together, these and a vast collection of other findings suggest that the mind is both structurally and functionally modular, to some extent.

However, both modular perspectives are weak in terms of their explanatory potential, as they tend to limit themselves to higher-level functions when they should be focusing on lower level components of a function to gain a clearer understanding. Furthermore, methodological biases and technological limitations present barriers to recognizing the cognitive versatility of the brain, i.e., we tend to think of brain regions as uni-functional, when in fact they might be multi-functional, so when brain activity is found to be associated with a certain function we stop looking for other functions that might activate the same area. As a consequence, a modular perspective facilitates the identification of specialized functions (e.g., language), rather than specialized operations (type of processing; algorithm).

For example, in the case of Broca's area, a modular perspective might cause a researcher to prematurely define it as a "language area" and direct all further efforts at trying to better understand its language producing capabilities, whereas a researcher employing a neural reuse perspective would look beyond language for the operation that is being carried out by the neural circuitry in relation to other tasks. The latter perspective clearly has a greater explanatory potential. In Chapter 3, I will attempt to focus on the FFA's operation (i.e., it's working; see below).

Using the massive modular approach as a reference point, the Fodorian module is clearly of only very limited use, as it does not provide insight into higher-level cognitive functions – the very functions we’re most interested in dissecting and for which we are currently making progress in understanding. The massive modularist’s conception of a module is more consistent with neuropsychological findings, however I will argue that neural reuse is more consistent than massive modularity). Indeed, massive modularity does not seem to stand up to the evidence coming from neuroscience and when examining neural activation across experiments (e.g., across domains, modalities and tasks), such as the findings associated with Broca’s area discussed in Section 2 (Bergeron, 2015). Together, the shifting back and forth between two modular perspectives, in addition to their weak ability to explain a wide variety of functions, creates an obstacle to understanding the actual cognitive architecture of the brain.

Neural reuse proposes that neural circuits established for a certain ‘working’ can be shared with other neural circuitry to support different uses, and are formed over both developmental and evolutionary timeframes (Anderson, 2010). The term ‘working’ can be thought of as a defined neural circuitry that performs a specific operation that can be reused to support multiple functional uses. A comprehensive understanding of how neural circuits are reused can provide profound insights into the brain’s cognitive architecture and avoid the constraints that are associated with a modular view. Indeed, focusing on “workings” or “working zones” can provide a more thorough explanation of function than attributing a particular function to a module.

Neural reuse has been proposed as a way to better understand cognitive function. Building on the work done by modularists, a neural reuse perspective can be an effective

tool in reassessing what can be salvaged (groups of specialized neurons) from modularity and what can be let go (encapsulation and domain specificity), or at the very least released from being a requirement. However, before diving full force into neural reuse, it might be advantageous to first explore the principles of its building blocks, by examining Bechtel's 'mental mechanism'.

Section 7: The Mental Mechanism – A Precursor to Neural Reuse

William Bechtel has helped develop a mechanistic approach to understanding the mind/brain. This approach is concerned with the mechanisms that produce mental phenomena and, more precisely, the operations and tasks associated with these mechanisms. I believe Bechtel's explanation of mechanisms facilitates a discussion on neural reuse as it concisely and independently does some of the work required for developing a rationale for a neural reuse perspective. In a few words, Bechtel's work on mechanisms highlights the importance of understanding the component parts within a system, both by way of their interaction with other components and the role they play in the overall system. There are obvious parallels here with respect to workings, their operations, and their role in supporting multiple cognitive uses.

Every scientific domain requires a tactic by which a relationship between two or more elements in a given system can be explained. Physics, for example, depends on laws (e.g., of motion, of gravity, of thermodynamics) to explain the nature of mass and energy. Other domains, such as psychology, depend on abstraction and the creation of concepts like memory or emotion to explain mental states. Biology, and chemistry to a certain extent, rely on mechanisms at various levels (i.e., molecules, cells, organisms) to explain how living things behave and function. Bechtel has advocated that the same

tactics used in biology can be applied to our examination of cognitive science (Bechtel, 2008). In other words, the study of mechanisms at various levels within the brain can help advance our understanding of cognitive functions.

Bechtel (2008) describes a mechanism as “a structure performing a function in virtue of its component parts, component operations, and their organization. The orchestrated functioning of the mechanism is responsible for one or more phenomena.” From this, we are to understand that the goal of the mechanistic approach is to decompose the mechanism into its parts and operations, and by doing so we will learn more about a given phenomenon. Bechtel (2008) explains that decomposing a mechanism is an iterative process, as often we discover that parts within the mechanism can be broken down further into sub-parts, and so on. This iterative nature of mechanisms gives rise to a layered configuration, consisting of levels of organization. Here the word “level” should not be confused with its other use, that which differentiates the scientific domain hierarchy that we are familiar with which begins with the more formally derived sciences and ascends to more conceptual sciences (i.e., physics → chemistry → biology → psychology).

Bechtel (2008) underscores that the operations at one grain size below the phenomenon are particularly useful in explaining the phenomenon; finer grains can shed light on coarser grain operations but cannot explain the phenomenon. To clarify, we can look to the gene as an analogy. At its most basic level a gene is made up of a strand of two paired bases: adenine(A)-thymine(T) and cytosine(C)-guanine(G), which are comprised of unique chemistries involving Hydrogen, Nitrogen, Oxygen and other elements. The sequence of ACTGs is what compose the gene and provide its function

(i.e., assembling a particular protein). To relate this back to Bechtel's levels within mechanisms, by understanding the chemistry of A, T, C, and G, we still do not have enough information to predict protein formation. We first need to know the base pairing rules, and further, we need to know the sequence of these base pairs. For this analogy, at the base level, the elements represent the parts and the chemistry that determines their configuration represents the operation. At the next level, A,T,C and G represent the parts and their pairing rules represent the operation, i.e., A always combines with T, and C always combines with G. Finally, at the top level, the sequence of base pairs represents the parts and their translation to a corresponding amino acid sequence represents the operation. The final product, the protein, is the phenomenon.

I have chosen to borrow from the field of genetics for this analogy because this mechanism has been well established. The same cannot yet be said for the field of cognitive science. It is important to see that how the parts at the level below the phenomenon are organized is what gives rise to the phenomenon itself. Bechtel (2008) describes this more or less in Gestalt terms, where the whole is greater than the sum of its parts.

In deconstructing a mechanism we learn that certain key parts perform operations. Bechtel (2008) proposes that it is necessary to make the distinction between “parts understood structurally from operations understood functionally.” It is important to make this distinction because different techniques are used to explore these two properties (Bechtel, 2008). Furthermore, if a researcher is fortunate enough to identify the properties of one, it is unlikely, or would at least be a challenge, to identify the properties of the other (Bechtel, 2008). Consequently, it is often difficult to associate parts with

operations (Bechtel, 2008). As explained earlier in this paper, Bechtel (2003) has advocated for a *components of processing approach* to examining and explaining function. This approach effectively brings the idea of a mechanism down to a much smaller grain size than that typically described by Fodorian or massive modularity. It is possible that some of Bechtel's thinking in this regard has been influenced by the work of Deacon (1997). Indeed, Bechtel draws significantly from Deacon's arguments against narrow evolutionary psychology. Bechtel's work on the mechanism is important because it forces us to imagine modules below the surface, and even below the surface of their parts. This enables one to understand that the parts of a system provide their own functions and that these functions might be so rudimentary that they can be exploited for many different purposes. The topic of granularity will be covered in more detail in Section 11.

I hope I have provided a sufficient overview of what Bechtel means by mechanism and how his approach of decomposing mechanisms can be useful for dissecting and understanding the mind/brain. Ultimately, it could be said that Bechtel's *mechanism* forces us to look below the surface of functions. Once the lens is directed at the inner workings of a module, neural reuse can take us to the next stage of understanding – the explanation of how neural circuitry (working zones – the neural anatomy that makes up the working and carries out the operation) can be co-opted by other groups to produce multiple functions, and how modules can appear to be specialized to perform a certain function when in fact we are seeing just one use of a more widely applicable working.

Section 8: General Theories of Neural Reuse

It might not be entirely true that Bechtel set the groundwork for the neural reuse perspective, but it occurs to me that his work has inspired it by pointing us in the right direction. By analogy, his work on biological mechanisms has provided interesting insights into how we might apply our understanding of biology to advance our understating of cognitive architecture.

Decades before Bechtel developed his work on the ‘mechanism’, Aleksandr Luria presciently described a framework for understanding cognitive function in his 1973 book, *The Working Brain: An Introduction to Neuropsychology*. The general claims he puts forward in this book share a number of similarities with neural reuse perspective. In particular he describes a non-localized cognitive architecture; a system comprised of component parts:

It is obvious that the whole of this process is carried out, not as a simple ‘function’, but as a complete functional system, embodying many components belonging to different levels of the secretory, motor and nervous apparatus. Such a functional system, differs not only in the complexity of its structure, but also in the mobility of its component parts (p. 27)...Naturally all mental processes such as perception and memorizing, gnosis and praxis, speech and thinking, writing, reading, arithmetic, cannot be regarded as isolated or even indivisible ‘faculties’, which can be presumed to be the direct ‘function’ of limited cell groups or to be ‘localized’ in particular areas of the brain (p. 29).

The components of Luria’s functional system resembles the “working” used by Bergeron (2007) and Anderson (2010) to describe the neuroanatomy that is active in supporting multiple uses within a neural reuse framework. In fact, Luria (1973) uses the same term in dissecting his ‘functional system’ into parts:

That is why mental functions, as complex functional systems, cannot be localized in narrow zones of the cortex or in isolated cell groups, but must be organized in

systems of concertedly working zones, each of which performs its role in complex functional systems, and which may be located in completely different and often far distant areas of the brain (p. 31).

Luria (1973) also notes that the use of double dissociations, a useful technique to identify supposed modules, is prone to error when it comes to localization as disrupting a single brain region likely leads to a collapse of the entire functional system and thus tells us nothing about the localization of a particular function. Further similarities are found in the sharing of parts, or workings, between functional systems. Indeed, the way in which Luria describes this sharing of ‘factors’ exhibits obvious parallels with neural reuse, in particular, when considered against what I have described of Broca’s area when neural reuse perspective is applied to this area:

Would anybody be prepared to accept at once that such different psychological processes as spatial orientation, arithmetical calculations, and the understanding of complex grammatical structures have important links in common, on the basis of which they can be united into a single group of psychological processes? As I have already mentioned, a lesion of the left parieto-occipital region of the cortex almost invariably leads to a disturbance of all these processes, so that a patient with such a lesion not only finds it difficult to find his bearings in space but also makes mistakes in his simplest calculations and misunderstands complex logico-grammatical structures. This means that all these apparently so widely different functions incorporate a common factor, and it allows an approach to be made to the more intimate analysis of the structure of psychological processes (p.41-42).

In addition to the work of Bechtel and Luria, a number of researchers have observed the importance of using different forms of neural reuse as an explanatory tool for understanding elements of the brain. Most of these observations have revealed themselves through specific case studies relating to some aspect of cognition, but Anderson (2010 & 2014) has thought about this topic more broadly and has developed a framework to support a global view of the brain.

Anderson (2010) highlights four general theories of neural reuse that have been developed in recent years: neural exploitation; shared circuits model; neural recycling theory; and massive redeployment. It is worth briefly examining each of these theories as each provides a distinctive view of the same general concept, however Anderson's favourite, massive redeployment, is the approach I will use to examine the fusiform face area in Chapter 3, as it is the most well-developed of the four. In other words, it has significant explanatory potential.

Neural Exploitation Hypothesis

The neural exploitation hypothesis proposes that various elements of social cognition derive from neural mechanisms that originally evolved for the purpose of social integration (Gallese, 2008). Gallese (2008) explains that the same premotor circuitry that is responsible for controlling action execution and the simulation of actions observed in others also plays a critical role in the hierarchical ordering of language and thought. It has been suggested that the two functions are supported by a common sub-level operation, 'phase structuring' which stimulate efferent neurons in the proper order (Gallese, 2008). The effector regions for these two examples are functions like action execution, imitation and imagination for the former; and formation of sentences or chains of thought for the latter (Gallese, 2008). Gallese (2008) admits that the idea of reuse of the premotor neural circuitry might disappear with more advanced imaging technology with greater acuity, however insists that this line of thought should be considered further.

Shared Circuits Model

In explaining her shared circuits model (SCM), Hurley (2008) proposes that the neural mechanisms responsible for perceiving action and for situated social cognition are created from active perception mechanisms. Hurley (2008) describes SCM's central hypothesis as "associations underwriting predictive simulation of effects of an agent's own movement, for instrumental control functions, can also yield mirroring and "reverse" simulation of similar perceived movements by others."

Effectively, perception and action share common coding. Certain neurons, *canonical neurons and mirror neurons*, have been identified as candidates for linking perception and action processing (Hurley, 2008). Canonical neurons "fire when an animal perceives an object that affords a certain type of action and when the animal performs the afforded action," whereas, mirror neurons "fire when an animal perceives another agent performing a type of action, and also when the animal performs that type of action itself" (Hurley, 2008). These subsets of neurons have been particularly useful in attempting to explain neural mechanisms for human imitation. Thus, Hurley's SCM involves the integration of input and output systems, where a subset of neurons are responsible for the shared resource that brings the two systems together.

This example does not fit perfectly with Anderson's more developed theory of neural reuse, as the working arguably supports the same use, or at minimum a very similar use: the imagination and carrying out of a particular action. However, it does reveal an interesting flexibility of neural circuitry and provides a point from which one could imagine this principle being applied broadly for hypothetical circuitry that supports multiple, mutually exclusive and functionally distinct uses.

Neuronal Recycling Theory

Dehaene (2005) has proposed that the brain contains domain specific mechanisms for language and arithmetic that originally evolved for other related but distinct functions. Dehaene (2005) points to a plethora of neuropsychological (lesion studies) and neuroimaging evidence to support this claim. In particular, this evidence reveals that a horizontal segment of the intraparietal sulcus is active solely during calculation (Dehaene, 2005). Research on various animals provides hints about past functions that the calculation mechanism might have previously supported in our historical lineage, and from which we co-opted this area to apply to arithmetic. For example, rats, pigeons and monkeys can extract the approximate numerosity of auditory or visual sets of objects (Nieder & Miller, 2004). Indeed, activity of neurons in the dorsolateral prefrontal cortex and intraparietal sulcus was recently identified in macaque monkeys in response to numerosity tasks (Nieder & Miller, 2004). It is thought that this localization in macaques is homologous to the human neural anatomy responsible for arithmetic (Dehaene, 2005).

While the majority of human neuroimaging studies on this topic have examined symbolic stimuli rather than numerosity, Piazza et al (2004) designed an experiment to habituate human participants, through repetitive exposure, to a specific visual numerosity (i.e., 16 dots) with the subsequent introduction of a single deviant numerosity (range from 8 to 32 dots). They found that the only brain region to respond to this change was the horizontal segment of left and right intraparietal sulcus (Dehaene, 2005). Moreover, In terms of neuropsychological findings, lesions of the horizontal segment of the intraparietal sulcus during development are associated with dyscalculia (Dehaene, 2005). Interestingly, this area has also been linked to the representation of social status hierarchy

(Chiao et al, 2009), which shares the characteristic of magnitude assessment in common with the previous examples.

Based on these and other findings related to language development, Dahanne (2005) has proposed the neuronal recycling hypothesis, which “emphasizes that cultural acquisitions must take place within the limited surface and bounded plasticity of the human cortex.” In this vein, Dahanne (2005) claims that evidence from neuropsychology and fMRI experiments relating to language and arithmetic suggest a higher than expected amount of reproducibility associated with these functions. In other words, it is conceivable that evolutionary neuro-anatomical and functional precursors play a significant role in setting the boundaries around culturally motivated behaviours. For humans, it is conceivable that neural circuits were selected for assessing and issuing social status hierarchies, for example, were later redeployed to develop and perform mathematics and language.

Massive Redeployment Hypothesis

The massive redeployment hypothesis has similarities with the three previously described theories of neural reuse, both in terms of maximizing efficiency and making use of neural plasticity. It also builds on work by Bergeron (2008) in which he employs the concept of a *working*. Building on my brief mention of this term at the beginning of this chapter, Bergeron’s *working* is described as “any bound region of the brain to which a cognitive ‘function’ can be attributed. Such brain regions can vary widely in size and shape depending on the particular nature and purpose of the proposed functional specification. Importantly, cognitive working zones may or may not correspond to recognized anatomical structures (e.g., Broadman areas, gyri, nuclei).”

Bergeron (2008) further defines a cognitive ‘working’ as the operation that a working zone performs. A working is a multi-domain operator that contributes to functions across perceptual modalities and other cognitive tasks (Anderson, 2014). Importantly, Bergeron (2008) underscores that “cognitive working zones remain invariant across very different cognitive contexts.”

To further clarify what he means, Bergeron provides the example of Broca’s area as a cognitive working zone. We can hypothesize that the cognitive *working* is that of a hypersequential processor that can be used, in collaboration with other brain areas, to produce a number of functions, including those associated with language, music, object manipulation and action perception (Bergeron, 2008). Evidence for this functional interpretation of Broca’s area largely comes from the work of Russell Poldrack (2006) who showed, through Bayesian analysis, that this area is more often activated by non-language tasks than language ones (Anderson, 2014). Bergeron’s description of working zones and their workings lays a foundation for Anderson’s conceptual framework of cognitive function by neural reuse.

At a high level, Anderson (2010) explains that evolutionary processes have favoured the reuse of existing neural mechanisms for new functions (I will return to the topic of neural reuse and evolution for a full discussion in Section 9). Anderson (2010) notes that according to this assumption we would expect that: 1) brain regions tend to support multiple functions; 2) older brain areas should be more widely reused, i.e., play a role in a wider variety of functions than newly developed brain regions; and 3) newer functions should use a broader range of brain regions dispersed across a larger area.

Anderson (2010) reports on a variety of research findings that support his three expectations. First, a review of fMRI data from almost 1,500 experiments reveals that a typical cortical region is activated by tasks in nine different domains (Anderson, 2010). Notably, domains were relatively incongruent in functions, ranging from vision to memory to attention and mathematics. These results were replicated when dividing the brain into significantly smaller regions (1,000 in total) (Anderson, 2010).

Regarding Anderson's second expectation, that 'older brain regions should be reused more', he points to research that demonstrates that the number of tasks in which a brain region is active increases as one moves from the frontal regions to the back (Anderson, 2010). Thus Anderson's expectation holds, as it is consistent with our general understanding that the back of the brain is older than the front (Anderson, 2010). Finally, Anderson (2010) reports findings that showed that among several domains tested (including reasoning, memory, emotion, mental imagery, perception, action, and attention), the neural circuitry for language, a relatively newly developed function, was the most widely dispersed throughout the brain. These findings present evidence in support of Anderson's third expectation: 'newer functions should use a broader range of brain regions'. I will return to the topic of evolution and neural reuse in Section 9.

The foundation upon which Anderson's massive redeployment hypothesis is based, is the proposition that "local circuits may have low-level computational workings that can be put to many different higher-level cognitive uses." Accordingly, Anderson (2010) asserts that cognitive functions should be understood by the unique way in which they combine shared workings. A co-activation analysis conducted by Anderson (2010)

revealed that every task was characterized by co-activation between regions, contrary to what would be expected from a modularity perspective.

Moreover, results from similar experiments demonstrate that the relative amount of brain region overlap between domains and functional connections between brain regions, across ten cognitive domains, was as expected under a neural reuse organization i.e., high brain region overlap and low connection overlap (Anderson & Penner-Wilger, 2012). In other words, the results revealed that the range of ten domains examined share brain regions, but these shared regions are specialized for a suite of uses. The alternative modular or holistic organizations would be represented by low brain area and low connection overlap, and high brain region and high connection overlap, respectively (Anderson & Penner-Wilger, 2012).

A cognitive architecture that involves the sharing of neural circuitry to support multiple functions would obviously seem susceptible to a sort of functional interference when a given working is employed for multiple tasks at a particular moment in time. Indeed, one would be tempted to assume that if a working is called upon twice at the same time, one or both workings will exhibit a decrease in performance or perhaps a loss of function all together. Anderson (2014) provides a number of examples of this functional interference.

First, Anderson (2014) references research by Glenberg and Kaschak (2002) which asked individuals to determine if a sentence presented to them made sense by pressing a button that was either close to their body or away from it ('yes' near body and 'no' away, for example). They found that when the response required a motion that was opposite to the movement expressed in the sentence (e.g., when a response to "put a

grape in your mouth” required a button push away from the body, or “you gave him the paper” required a button close to the body), it took more time for the participant to respond. It is thought that this delay reflects a shared neural circuit that is being used to support two tasks: motor control and language processing (Anderson, 2014). Other research suggests that neural circuits in the left precentral gyrus and left angular gyrus are shared for the uses of finger awareness and number representation (Anderson, 2014). In this case, circuitry evolved for the former was likely co-opted for the latter (Anderson, 2014). In addition to these examples from the domains of language and numerical cognition, Anderson (2014) also provides evidence for neural reuse in memory structures, spatial cognition and perceptual structures.

Given its incompatibility with a strong modularity perspective, Anderson (2010) draws from global wiring optimization theory to position the role of neural reuse on the overall functioning of the brain. Wiring optimization theory explains that the length of connections between regions of the brain is minimized, however this does not mean that areas that experience co-activation are adjacent to each other (Anderson, 2010). In fact, evidence shows that more recently acquired cognitive functions are typically farther apart in the brain (Anderson, 2010). It seems then that the evolution of function is involved in a constant tug of war between energy efficiency and increased fitness provided by a particular function. It is important to compare the evolutionary pressures supporting energy efficiency and adaptability benefits derived from some form of modular organization, as described by Clune et al (2012) (in above section: Evolution, Development and Learning), to the efficiency and functional advantage of building on what’s already there.

Section 9: Neural Reuse Framed by Homology and Development

As discussed in Section 5, any theory of mind must be compatible with evolutionary theory. The objective of this section is to demonstrate how the neural reuse perspective aligns with what we know about the evolution of other species, and the development and plasticity of the human brain. I also intend to show that our knowledge of these areas aligns more with a neural reuse perspective than modularity. While it's difficult to surmise whether cognitive architecture derives more substantively from developmental or evolutionary roots, Anderson & Penner-Wilger (2013) are confident that the neural reuse perspective offers opportunities for new insights relating to both accounts. In terms of development, the neural reuse perspective has its roots in embodied cognitive science (ECS), which places significant emphasis on interactions between neural structures and environments, and focuses on efficiencies of the whole brain rather than efficiencies of its component parts (Anderson, 2014). Indeed, ECS shares with neural reuse the assertion that evolution should favour the reuse and redeployment of existing parts, rather than the creating a solution from scratch, in each case, as massive modularity requires.

Anderson & Penner-Wilger (2012) underscore that on the evolutionary front, we should observe structural neural homologies across species that carry out different functions. For example, the homologue of Broca's area in primates should support a different use than that of language by employing the same working. This could provide insight into the working of this brain area in humans. Similarly, Anderson & Penner-Wilger (2013) note that we would expect that two functions arising at the same time in development might both make use of the same neural circuitry that developed earlier for

a distinct function, if the operation that is carried out by this circuitry is useful and applicable to each (Anderson & Penner-Wilger, 2013). Clearly, both play a role in setting the parameters for neural reuse, however it has not yet been determined with certainty what that role is and how significant a role each performs.

In response to Anderson's 2010 paper, Bergeron (2010) provided additional insights into 'functional homology'. Bergeron (2010) highlights the ambiguity that arises with the term 'functional homology' in the context of neural reuse, as it is the *working* that is conserved across species, not the high-level function (use). Likewise, it is the working that is conserved within species and within the individual as well. In 2003, William Bechtel presciently sketched a similar diagnostic for functional underpinnings to that of Bergeron. Bechtel (2003) has sought mechanistic explanations for high-level cognitive functions. In doing so, he has turned to evolutionary psychology; in particular, by searching for phylogenetic linkages across species by means of, what he describes as, "decomposition of the performance." Bechtel's decomposition appears to effectively be a dissection of function into Bergeron's workings, essentially a low-level micro-function that is understandable against an evolutionary backdrop, i.e., these workings are functional units that were actually evolved over the broad timeframe of our species and beyond.

Bergeron (2010) proposes that Broca's area is a good candidate for cognitive homology across species. In this case, the homologous working is "the detection, extraction and/or representation of regular, rule-based patterns in temporally extended events", which supports unique functions/uses across species e.g., action sequencing and manipulation of objects in primates, and language production in humans (Bergeron,

2010). There exist obvious similarities here with what we know about the intraparietal sulcus, e.g., low-level heuristics that contribute broadly to different functions.

In response to Anderson's views, described above, relating to the merits of homology, Katz (2010) opines that since neural development and organization is strongly based on inputs the brain receives, it is problematic to talk of homologous brain areas across species. As an example of this, Katz (2010) refers to the fact that the anterior portion of the visual cortex, which is normally dedicated to vision, can be employed for reading brail. Katz (2010) underscores that it is not the function (i.e., the working) that is preserved, given that a particular function depends on the antecedent inputs this brain region has received, rather it is the algorithm. There are obvious similarities here between Katz's use of algorithm and Anderson and Bergeron's use of workings. This algorithm could be thought of in the same way as Bergeron's working; it provides an operation that contributes to multiple functions, e.g., aspects of vision or reading brail. Understanding the role of these workings/algorithms and how they combine with others to produce function will be key to understanding cognition. In particular, understanding at which level input acts to mold the brain is paramount. It is conceivable that input acts much like a musical conductor, singling out certain workings and encouraging them to combine with others over time.

Dekker & Karmiloff-Smith (2010) further substantiate the propositions of Anderson & Penner-Wilger (2012) and Bergeron (2010) regarding consistencies between a reuse perspective and our understanding of brain development. Dekker & Karmiloff-Smith (2010) consider the impact on function associated with developmental disorders. Their contention is that the belief by many researchers that developmental disorders

typically result in the loss of a specific function, and thereby imply some sort of modularity, is wrong (Dekker & Karmiloff-Smith, 2010). They suggest that for developmental disorders associated with abnormal brain formation, subtle deficits in intact domains are continually overlooked (Dekker & Karmiloff-Smith, 2010). For example, Dekker & Karmiloff-Smith (2010) explain that Williams syndrome, which is caused by a decrease in gene activity on chromosome 7, is characterized by impaired activity across multiple cortical regions and deficits in numerous cognitive domains.

Based on these findings, Dekker & Karmiloff-Smith (2010) have determined that affected genes are unlikely to be precursors to specific modules. It is conceivable that these genetic impacts can lead to a disruption in the development of early neural circuits, resulting in a cascading effect where the working associated with this circuitry will have deleterious effects, ranging in magnitude, across a number of functions that would typically rely on the working.

A Developmental Explanation for Neural Reuse

To this point in Section 9, I have focused on the mechanics of development in general terms. However, in order to better demonstrate that neural reuse is a superior perspective to modularity an explanation of how actual workings develop is required. In this section, I will summarize recent thinking of Anderson (2014) as to how workings are formed. It should be noted that what the authors, discussed below, call ‘functional’ includes what Bergeron (2007) refers to as a working, which I have chosen as the preferred terminology throughout this paper. For consistency, this should be kept in mind while reading this section.

Anderson (2014) presents a comprehensive framework for how a brain is organized in a way that supports workings that are reused; Anderson calls this the Interactive Differentiation and Search (IDS) framework. In developing this framework, Anderson (2014) builds on earlier work by Johnson (2001, 2010) who has put forward an Interactive Specialization (IS) framework, which describes functional development in the brain as “changes in the response properties of cortical regions during ontogeny as regions interact and compete with each other to acquire their roles in new computational abilities” (Johnson 2001, p. 480).

Indeed, the IS framework argues that some cortical areas begin with a wide range of functional potentials but are narrowed down and become specialized as they compete with other areas over a particular function (Johnson, 2011). In addition, IS proposes that functional development is dependent on the interaction and organization of multiple brain regions (Johnson, 2011). Thus, learning new cognitive abilities requires a reorganization of existing networks, rather than an addition of a new dedicated region (Johnson, 2011).

In developing his IDS framework, Anderson (2014) transforms the “specialization” aspect of IS such that the process becomes one of differentiation, rather than specialization. This is to say that,

Over the course of development, local neural assemblies will come to have particular, distinctive response profiles, as determined by a combination of intrinsic local cortical biases and extrinsic factors including experience and the influence of functional interactions with other regions of the brain. A region’s response profile will certainly reflect its underlying functional capacities and determine the role(s) it can play in various functional coalitions. But although this might therefore be considered a kind of functional selectivity, it is quite different from the notion that brain regions will come to specialize in such tasks as “face perception” or “mind reading” and furthermore suggests a development pathway relatively unconstrained by (or simply in some sense insensitive to) the traditional categories of cognitive psychology – p. 52-53.

Anderson's IDS framework accepts the general IS claim that functional development is dependent on the interaction of brain regions, however he deviates from it in how a particular working is established for a given use and how flexible a working is to carry out other uses. Whereas IS describes a process in which neural circuits compete against one another (through the interaction with one another in a "learning environment") for a specific type of working and come to be specialized in carrying out one particular use that fits a given neural circuit best, Anderson (2014) proposes that neural circuits undergo a similar learning process regarding their own aptitudes and capabilities versus other circuits but in the end each circuit is assigned a working that can support multiple uses. One alternative to these views is that functional development is genetically predetermined and non-modifiable (i.e., a process guided entirely by nature). Clearly neural reuse is genetically driven, however its outcomes, i.e., its workings, are highly dependent on environmental inputs.

The "search" in Anderson's IDS framework reflects the brain's proposed ability to challenge a range of networks with novel sensory input and determine which is best positioned to do work with this input (i.e., learning). Anderson (2014) describes the challenge of learning as an attempt to "find (and fine-tune) the coalition of neural partners that have the right response tendencies – the neural architecture that has the right mesh with the input – and will therefore do something useful or adaptive in light of the sensory instructions."

In unpacking his claim that such a search actually occurs, Anderson (2014) proposes a mechanism to explain this process. Working from evidence suggesting that neurons can participate in more than one functional circuit, possess a certain level of

plasticity, i.e., can form new and different connections and can be influenced by neuromodulatory effects (e.g., volume transmission – the nonsynaptic diffusion of neurotransmission), Anderson (2014) proposes that during learning neural partnerships form for a period of time, by means of volume transmission, and those partnerships that are most effective become established anatomically. Anderson (2014) emphasizes that, at this point, we can only speculate about how far-reaching these actions can be (i.e., at what maximum distance volume transmission can have an effect in terms of establishing a working).

In an effort to further clarify the distinction between IS and IDS, Anderson (2014) compares how each framework would assess developmental prosopagnosia. An IS framework would explain this face perception deficit as a result of the working for this use being blocked or interfered with to prevent the working from being established; in other words, preventing specialization from occurring (Anderson, 2014). On the other hand, IDS would explain this type of prosopagnosia as a result of the search not being carried out properly (establishing a suboptimal working), not fully (e.g., it was interrupted), or not at all (Anderson, 2014).

There is not yet sufficient empirical evidence to determine the legitimacy of Anderson's IDS framework; Anderson concedes this himself. However, it is an interesting hypothesis that could provide an explanatory basis for preferring the neural reuse perspective to modular perspectives, as it provides an explanation for function rather than merely a description of it – it tells us how the brain employs a neural reuse architecture rather than simply providing a description of what neural reuse involves. A number of considerations first need to be fleshed out. In particular, the mechanics of

volume transmission need to be better understood. How a neuromodulator causes a configurational change to a working based on a given stimuli in some neural circuits but not others is currently unclear, and determining a mechanism by which this occurs could strengthen or weaken the validity of Anderson's IDS framework. Indeed, identifying properties of volume transmission and the neural circuits that lead to this selection bias (relating to the successful candidate circuit) would provide real legs to this hypothesis.

Currently, it's not clear whether the proposed volume transmission effect has something to do with the concentration of the neuromodulator being dispersed, or the type of modulator, or the number/type of receptors on the neurons that are interacting with the neuromodulator. Or perhaps it's a tug of war between more than one of these elements? It is also unclear at which level(s) of granularity this occurs. We should ask whether this framework also applies to sub-workings (i.e., the component neural parts of workings), the sub-workings of sub-workings and so on? It is important that further research be undertaken to clarify these questions.

Section 10: Are Workings Fixed?

I have dedicated this section to examining whether workings are fixed, and if they are, how they might become fixed, and at what point they become fixed (e.g., at birth, some point during development or in maturity). Providing clarity on this issue can do three things: lead to a more detailed description of neural reuse, inform our assessment of Anderson's IDS framework, and further fill out the incompatibility between neural reuse and massive modularity. However, a first step is to determine if there is evidence to support a fixed workings approach.

Donnarumma et al. (2010) provide insight on whether a fixed system can be implemented in practice. They provide a “fixed-weight programmable networks” model to explain how a programmable neural network can be developed and mimic the action of reusable workings in real time. Second, Dekker & Karmiloff-Smith (2010) cite evidence in support of fixed neural workings, and for a flexible, environmentally directed establishment of these workings. Dekker & Karmiloff-Smith (2010) suggest that this apparent conflict can be explained from a neuroconstructivist approach which could allow for early development giving rise to coarsely coded areas that are better positioned for certain functions but are also exposed to competition from other areas. Together, this evidence points to a process by which workings are fixed over the course of some developmental time frame, and thus a process that is genetically predisposed rather than predetermined.

Dekker & Karmiloff-Smith (2010) assert that fixed neural circuitry for workings would mean that tasks previously supported by a combination of workings would not be altered if one or several of the workings were annexed to support a new function. However, it is important to consider which grain size we are talking about when we talk of a working being fixed. The concept of a working has not yet been well articulated regarding at which grain size of the brain it resides; generally speaking, it is somewhere between the neuron and a lobe, and depending on which working is being considered, it can be found in any of these levels. Thus, workings can be comprised of sub-workings that can also be reused. It is for this reason that a working (or sub-working) can be located within the broad spectrum between lobe and neuron referred to above. Based on this building block model, changing a working would seemingly lead to a disruptive

snowball effect, presumably causing a disruption to some, if not all, uses; as a mutation to our DNA most often leads to negative effects, it would seem that altering a working would have the same tendency to negatively impact most of its uses. It follows that if components of workings were susceptible to change then this snowball effect would be amplified further, as it relates to functions. If a working is to be fixed, must its component parts (sub-workings) also be fixed? As I will discuss below, in the context of neuromodulation, the significance of this concern depends on whether the change is permanent or temporary.

Before answering the question about whether a fixed workings' component parts must also be fixed, it is important to consider what a 'working' in this context means. As indicated above, a 'working' can be thought of in two ways: as a system of parts that has a number of uses (single system working); or a component that combines with other workings to support multiple uses (component of a system working). While this distinction represents two sides of the same coin, this dual understanding of what a working looks like can cause confusion at times, as it can become unclear if a working's structure is being transformed for each use it supports, or if it is simply collaborating with other workings to support different uses.

With respect to a 'component of a system working', the legitimacy of neural reuse is based on the idea that a working can form partnerships with a variety of other workings, where each partnership can support a unique system use. As discussed in previous sections, partnerships are thought to be formed by learning and experience and in some cases can be predisposed by genetic factors. If one digs down within a working, it can be broken into component parts; again these parts are free to be reused with the

parts of other workings to support their own workings. Based on this description of workings and their sub-workings (and so on) we can consider a working to be fixed, in that its suite of neurons remain the same, but the multi-systems (multiple partnerships) that these fixed workings serve can be considered dynamic.

A concern raised by Dekker & Karmiloff-Smith (2010) is that Anderson is not clear on whether formation of new partnerships among neural workings is the only way that neural reuse can lead to new uses. Dekker & Karmiloff-Smith (2010) refer to Dohaene's neural recycling hypothesis as it allows for localized plasticity as a means to acquiring new uses such as mathematics and reading. Under this scenario, functions that were previously supported by the neural circuitry that was coopted and reformed to support the new use would be affected (Dekker & Karmiloff-Smith, 2010). Immordino-Yang et al (2010) present a cultural analysis that supports this consideration; they propose that cultural differences can lead to distinct responses that involve the altering of workings that have been ontogenetically or evolutionarily formed for another use. For example, responses to psychosocial distress of two populations (i.e., somatically for some Asian populations versus depression for westerners) are thought to be influenced by cultural differences in the recruitment of neural systems for somatosensation (Immordino-Yang et al, 2010). What is important to remember here is that workings that support uses beyond that for which they were originally selected do not require a structural change, i.e., they can be fixed. But, what if one or more components (sub-workings) of a working were altered, by development or evolutionary pressures?

Anderson & Penner-Wilger (2012) attempt to shed light on the concern that changing a working will have disruptive downstream effects (by effecting all uses that

share the same working) in a way similar to that of Dekker & Karmiloff-Smith (2010). They acknowledge that the more ways a working is shared in support of different uses, the harder it will be to change this working without impacting the overall viability of the system, as a change to a working that creates a benefit for one use might also be associated with deficits to other uses (Anderson & Penner-Wilger, 2012). Based on this assumption Anderson & Penner-Wilger (2012) conclude that neural reuse should constrain the flexibility of development. As evidence for this conclusion Anderson & Penner-Wilger (2012) note that learning a task early in life can make it easier or harder to learn a subsequent task, later in life. For example, pitch training can predispose children to better mathematics results later in their education (Anderson & Penner-Wilger, 2012).

Now, with respect to the first understanding of a working, i.e., a single system working, Anderson (2014) proposes that a working's conformation can be changed by neuromodulators, resulting in a new novel operation supporting a new use (as mentioned above, this appears to be a deviation from his earlier work). In this case, the working is not merely forming new partnerships in support of a new use (an easy process to conceptualize), rather the working itself is being changed in some way to support a new use. Indeed, the working is held in a state that can support a set of uses under certain neuromodulatory conditions, or it can be held in a new state (or multiple new states) that support a separate set of uses under different neuromodulatory conditions (Anderson, 2014). Although it is not entirely clear how Anderson understands this to happen, perhaps it can be assumed that the same network of neurons that the working is composed of are fixed but acting differently together, leading to the performance of a new operation.

Further to this, Katz (2010) suggests that neuromodulatory effects influence the multi-functionality of neural circuits. However, in the context of neural workings, it is not clear whether this neuromodulation is theorized to represent a functional change to the working, a physical change to the working, or a functional change to a collection of workings wherein the physical properties of each working remains constant. Clearly, if a neuromodulatory effect is one that simply ramps up or turns down the operation (e.g., analogous to increasing or decreasing the flow of water from a tap) while leaving the type of working intact (e.g., the output, water for example, does not change, to wine for example), there exist no problems here for neural reuse theory.

However, if neuromodulation in some way changes the working, these scenarios could present problems for elements of the neural reuse perspective as it has been sketched thus far, in that a neuromodulatory-induced change to a working would imply that it is not the working that is conserved to perform multiple uses but that the characteristics of a neural circuitry are changed to produce multiple workings that each support a unique use. Indeed, this seems to be what Anderson (2014) is proposing, which is a significant deviation from his past work (e.g., Anderson, 2010).

There are a number of ways that this neuromodulatory effect could be understood. First, if neuromodulation changes the function of a working, this adds a whole new level of complexity and significantly widens the flexibility with which workings operate, a characteristic that fits well within an evolutionary context. Second, if neuromodulation causes a physical change to a working, this presents a problem as it means that workings are not fixed, as proposed by Anderson (and further elaborated below). If workings are not fixed, one might presume that this would have a deleterious impact across the range

of uses to which they contribute, i.e., it is unlikely that a change to a working that provides a benefit to one use would provide a benefit to the other uses as well. However, if the neuromodulatory effect has a relatively short temporal effect, this concern might not be warranted. Third, a change in function derived from a collection of workings would be difficult to understand as this general modulation would likely spill over to other functions that share one or more workings, thereby resulting in the same snowball effect as with a physical change to a single working. Again, this concern can be dismissed if the neuromodulation is relatively fast acting.

Katz (2010) provides an example of neuromodulatory effects on function: The difference between pair-bond after mating in two vole species (prairie versus meadow) has been linked to regulation of the vasopressin V1, a receptor in the ventral pallidum. The V1a is abundant in prairie voles, which generally demonstrate pair-bonds after mating. Meadow voles, which do not pair-bond after mating, demonstrate increased pair-bonding following gene expression techniques that cause elevated expression of V1a receptors.

If Anderson's and Katz's hypotheses are true, we are left with a working that has a fixed structure but an operation that is not fixed. As discussed in the previous section, this does not necessarily present a problem for neural reuse but it is a deviation from past understandings of it. Indeed, it introduces another level of complexity in that a fixed neural circuit can be associated with multiple workings rather than a single one. In attributing this new multiple operation aspect to 'workings', Anderson has eliminated the problem introduced in the previous paragraph relating to the downstream consequences of altering a working, as it enables a working to wear multiple hats and therefore need not

introduce downstream effects. As this is a deviation from Anderson's prior description of a neural reuse (see Anderson, 2010), it may be worthwhile to review Anderson's historic views on whether workings are fixed in order to determine if they are compatible with his more recent views.

Anderson & Penner-Wilger (2012) have noted that neural reuse can be described as a developmental serial homology, in which the working is reused rather than being physically duplicated, which is most commonly the case in the context of homologies. Anderson & Penner-Wilger (2012) explain this as "the augmentation of overall function via temporal duplication of function." The implication here is that workings appear to be fixed, but is this consistent with Anderson (2014)? The answer to this question depends on the effect that neuromodulation is thought to have on a working. In terms of its basic structure, there is no reason to suggest that workings cannot be fixed, as the same neurons are present. However, if one looks beyond the neuron as a basic unit, the working might not be fixed in terms of the number of receptors found on a neuron and the number of dendritic connections. However, it is not clear at what point in development Anderson believes workings become fixed. If there is room for development to mold workings prior to being fixed, or almost fixed, then there is no reason why neuromodulatory effects on a working should be considered problematic or present a barrier to better understanding the brain by way of neural reuse. Rather it would simply provide an explanation of how reusable working are developed. Indeed, it may be the case that Anderson's serial homology is not in fact fixed, or fixed at birth. This will be explored further at the end of this section.

Clearly, there is a point at which it is likely that some level of neural circuitry is fixed – those workings that are old, in evolutionary time frames, and support a large number of higher level workings. Alternatively, it seems we would be left with a conundrum if they do not become fixed through development, as changes to a working would negatively impact a number of functions across the brain. In this light, the findings of Immordino-Yang (2010), for example, might be better explained by a reorganization of workings rather than an alteration of one particular working. In this vein, it would be interesting to examine the impact of genetically linked neurological disorders that might involve the disruption or alteration in the proper formation of one or multiple workings. Indeed, based on the premise that workings are widely shared, we would expect that these individuals would demonstrate deficiencies in all uses that the disrupted working supports.

Anderson's workings do not directly produce a single, narrowly defined function; rather it's their combination with other neural circuits that yield a particular use. Of course this invites the question, are there sub-level workings that combine to form sub-level uses? Workings that support uses that are too primitive (too far removed from understandable behaviours) for our eye to easily interpret. The answer, as alluded to throughout this section, is likely yes, and once workings are better understood, these lower level operations would enable us to understand the mind even further. This opportunity to dig deeper, combined with the similarities between workings and modules, does not mean that the neural reuse hypothesis is less effective than if workings and modules showed greater diversity. Rather, it demonstrates precisely how this perspective can bring us closer to understanding cognition, and the mind at large. In addition, while

it is easier to conceptualize a fixed working as one moves toward lower grain sizes, it is unclear at this point where, and in what sense, workings become fixed.

Section 11: Granularity

As observed in the previous section, knowing at which grain size of neural activity to focus has been a challenge for researchers when attempting to link anatomical brain areas to a particular function. The same problem exists when considering modularity, where fodorian modularists are concerned with relatively low levels, e.g., regarding shallow processing, massive modularists broaden their net to include more complex, higher level processing that integrate various shallow processes. Unlike Fodor's 'limited access' element of a module, which I previously assessed as unimportant from a neural reuse perspective, the idea of 'activity levels' (for which Fodor highlights those that are shallow) is paramount to considering function in this frame.

While "level" is a relatively convenient and commonly used term to describe at which point in the cognitive/ neural chain of events we are interested in looking, as with "module" it carries with it a significant amount of baggage. Indeed, David Marr and Tomaso Poggio (1976) put forward a Tri-level hypothesis that proposes information processing occurs at three levels: computational, algorithmic/representational, physical. While this characterization seems to aid in our understanding of some cognitive functions, e.g., the stages of vision, it does not help us to explain many other functions, particularly higher cognitive functions. Moreover, Marr and Poggio's use of "level" is not what I mean by the term. Rather than a differentiation of the levels of information processing, I am interested in examining the levels by which one can describe structural components of the brain, from lobe down to individual neurons, and their component

parts. To avoid any confusion on this topic, I will use the term “granularity” or simply “grain size” from here on, as it provides a better conceptual image of the physical attributes that I’m considering, rather than processing attributes described by Marr and Poggio. As such, an example of a coarse grained structure would be closer to the lobe end of the spectrum, whereas fine grain refers to a relatively small brain region or neural circuit. Coarse and fine grain can also be used within a given context to differentiate between the whole and its parts, e.g., Broca’s area (coarse grain) and its sub-components (fine grain).

Identifying the level at which fundamental operations or workings takes place is the essential first step in any examination of neural reuse, as it enables one to establish a starting point from which connectivity analysis can determine how a working behaves in relation to other workings. While the identification and examination of a single working is useful for understanding a single use or the full range of uses associated with it, if we are interested in explaining the general cognitive architecture of the brain it is necessary to explore the relationship between workings and also the relationship of components (sub-workings) within workings. In other words, it is important to examine the brain across levels, where higher-levels include larger anatomical areas and contain a larger number of sub-workings.

In addition to the concerns raised by Dekker & Karmiloff-Smith (2010) in the previous section relating to the fixedness of workings, Junge & Dennett (2010) raise an interesting hypothesis. They have recognized that it’s conceivable that a finer grained view of the brain would show that a type of neuronal specialization is being masked by what appears to be reuse (Junge & Dennett, 2010). To spell this out, Junge & Dennett

(2010) suggest that what Anderson interprets as reuse vis-à-vis workings might actually be a specialized neuron, or small group of neurons, that share identical characteristics with their neighbours who provide the same service in support of other functions.

In other words, groups of certain types of neurons, located at a much finer grain size than typically considered, that provide a particular working would be recruited individually to form unique collaborations with other brain regions to perform different functions. This would thereby represent multiple uses of certain types of neurons, where each neuron would support a single use (Junge & Dennett, 2010). Without a means to examine the activity of neurons, or small neural groups, this architecture would be easily misinterpreted to be one of reuse. This appears to be a valid concern and one that will remain a mystery until more sensitive neuroimaging techniques are developed. It should however be noted that Junge & Dennett's "specialized neurons" hypothesis would appear to be an inefficient use of resources. It is difficult to imagine that evolution would select for extra neurons to provide a use that is not yet needed, which seems to be the implication unless the number of these neurons are increased with time as selective pressures present themselves. Indeed, the reuse of neurons is intuitively more efficient than duplication of neurons, as there is a production and maintenance cost (where the currency is energy) associated with any living cell. Again, the problem proposed by Junge & Dennett (2010) is one that highlights the importance of considering which grain size is most important when seeking to explain a particular cognitive function.

The Junge and Dennett challenge leads to a more general issue that has already introduced itself several times in this paper. The lens that one uses to conceptualize a working can be zoomed in or out. Let's take Broca's area as an example; it is thought

that the working of Broca's area is that of a hierarchical sequencer, where Broca's area is defined by activity seen in fMRI analysis. Assuming all activity seen on the fMRI does reflect the activity of this working and not background noise, it is conceivable that this working is comprised of sub-workings that contribute their own special uses to the larger working (hierarchical sequencer) and possibly other larger workings. So at one grain size we are looking at the hierarchical sequencer 'working', but at a smaller grain size we are looking at sub-workings, of which some support the hierarchical working and other workings. As discussed above, if one is interested in examining a single working, it might not matter at what grain size this working is truly being realized as the function is what's really important in this case, not the anatomy. However, if one is interested in acquiring a more comprehensive understanding of the brain's cognitive architecture at large, it would be useful to explore the sub-workings of this working to determine if they can be linked to other uses, or to explore the connectivity of the working with other workings to understand its role in broader brain function.

A way of imagining grain size within the brain is to think of a photocopier. The working of the photocopier is to produce copies. However, if you examine the photocopier at a finer grain, you'll notice that there are two sub-workings: one that makes a copy from a paper that is placed into the machine and one that makes a copy from an electronic copy of a document that is sent from your computer. Each sub-working contributes to the overall working of the photocopier but in slightly different ways. Now imagine moving to a coarse grained view. Imagine the office in which the photocopier is located. The working of the room is to produce work in a certain area; let's say philosophy. In the room, there is a bookshelf that provides a sub-working of holding

books. The books possess the sub-working of holding information. The photocopier provides the sub-working of printing information. The philosopher provides the sub-working of organizing the other workings in a meaningful way. So, from this point of view, the working of the photocopier is actually a sub-working of the working of the office.

In a similar vein, Anderson (2014) highlights that confusion can arise when examining neural circuits at different levels of granularity. A given neural circuit, or the working it carries out, can appear to exhibit reuse at one level of granularity, but multiuse at a larger grain (Anderson 2014). Here is how Anderson (2014) makes the distinction between reuse and multiuse and the effect of neuromodulation (which was explored in greater detail in Section 9):

It might seem that the role of neuromodulation should instead lead us to strongly distinguish neural reuse from neural multiuse, reserving “reuse” for the case when a single neural element (neuron or network) is reused in the same state for multiple purposes and “multiuse” for the case when the element moves into a different functional configuration. . . . Because there is reuse at multiple spatial scales, what is reuse at one level or organization can be multiuse at another. Using the same neurons in a different configuration or when modulated by genetic or chemical factors is reuse of neurons, but multiuse of the local network. Reuse of a local region that cooperates with different partners is reuse of the region but multiuse at the level of the larger-scale network (p. 36).

So, to again borrow the room and photocopier analogy, a professor from the biology department (which is located right next door to the philosophy department) might use the philosopher’s photocopier (budget cuts have limited the number of photocopiers so they must share). As we know, in the philosopher’s office, the photocopier can be reused in two ways; printing and photocopying. However, when the biologist uses the photocopier (to stand on to reach a book off a shelf), the working now supports a higher

level working that is not the working of a philosophy office, rather the working of a biology office; an example of multiuse. Thus, from the perspective of the biology office and philosophy office, the working of the photocopier exhibits multiuse, and in fact, the working appears to perform a new operation using the same system. As discussed in earlier sections, this is a deviation to Anderson's previous work (2010) and leads to many new questions, such as what is driving the release of neuromodulators that supposedly act on workings? And what types of conformational changes are the neuromodulators causing? However, if Anderson (2014) is simply saying that neuromodulation causes one or multiple components of the working to behave differently, and that this results in a new operation, this does not present a problem, however it obscures our understanding of what a working is because the working that has undergone a change due to neuromodulation seems to be better explained as an unique new working rather than a configurational change to the same working. While the same neural circuitry is employed, the operation has changed and therefore the work it is carrying out has changed.

Section 12: At Odds With Massive Modularity

It can be said that neural reuse has merely replaced one form of modularity with another that is more flexible and low-level (Dekker & Karmiloff-Smith, 2010). Dekker & Karmiloff (2010), for example, are concerned that workings could conceivably be thought of as modules, if we stretch the parameters of what a module is to the extent that some massive modularists do. Indeed, as discussed in Section 9, Ritchie & Carruthers (2010) have declared that a module need not be encapsulated or domain specific. So, to what remaining characteristics do Ritchie & Carruthers (2010) turn in order to justify

their use of the term “module”? They maintain that their modules are dissociable and separately modifiable.

First, on the topic of functional dissociability, Bergeron (2007) highlights an important point regarding the distinction between two functionally independent processes; they need not be qualitatively different. What’s required is that they simply run independent of each other, i.e., one can be damaged without affecting the other (Bergeron, 2007). Moreover, Bergeron (2015) explains that while the view that double dissociations support modularity is an inference to the best explanation, these same outcomes can be observed for non-dissociable systems. Bergeron (2015) unpacks this claim using workings and the uses of workings to demonstrate that a use can be eliminated while leaving the working intact.

Bergeron uses the example of Broca’s area to add clarity to this; the hierarchical sequencer that is thought to be the working of Broca’s area can remain intact even though one of the uses it supports, e.g., language production, is damaged. In this case, some of the brain regions that the working collaborates with to produce language have been damaged, but the areas that it collaborates with to perform object manipulation, for example, remain intact (Bergeron, 2015). Thus while Broca’s and Wernike’s area are said to be functionally independent from a modular perspective, the same functional independence could be observed by applying a neural reuse perspective to Broca’s area, where independence is attributed to the reuse of a single brain region for multiple uses, rather than the dissociation of two brain regions. This nuance is not meant to speak to whether Ritchie and Carruthers’s description of a module makes sense, rather that the

evidence that is likely driving their decision to hold onto functional dissociability is not as compelling as they might propose.

With respect to Ritchie and Carruthers's definition of a module, Anderson (2010) raises an obvious objection: if modules can share parts, they cannot be functionally dissociable and separately modifiable. In response to this objection, Ritchie & Carruthers (2010) claim that although modules A and B share a common operation they also could contain neural circuitry that does not overlap. If this non-overlapping circuitry is damaged, they are dissociable and if this non-overlapping circuitry is altered they can be separately modifiable. As noted by Anderson (2010), this may be true but when one takes a broader look a significant problem becomes clear: if damaging non-overlapping circuitry in module A does not impact the function of B, it will almost certainly impact the function of some other module C, according to neural reuse (of course, none of these sites would be referred to as a module from a neural reuse perspective but using this terminology here is useful to make the point).

Taking this broader look, it becomes clear that neither functional dissociability, nor separate modifiability, could generally be the case according to the neural reuse perspective. While either can be present in unusually isolated areas of the brain, they could not apply to areas where the vast connectivity of the brain is in operation. Thus, even a massive modularist view as flexible as Ritchie and Carruthers's – one that does not require encapsulation or domain specificity – is not compatible with neural reuse. Carruthers's massive modularity would need to do one more thing if it is to align itself with neural reuse: drop the requirement that modules be functionally dissociable and separately modifiable. If Carruthers were to give this up, his 'module' would look very

much like a working. As claimed in the first part of this chapter, and which I hope has been supported throughout, neural reuse is a superior perspective than modularity for explaining the brain's cognitive architecture as it fits the facts better. I also believe that it has a greater ability to generate interesting predictions, and a stronger propensity to generate interesting questions, but that will be a question for further work.

In the following Chapter, I will attempt to use the neural reuse perspective to more precisely define the working of the FFA and to explore whether it is more versatile than simply a domain-specific device for recognizing human faces. I have selected the FFA because it has long been thought of as a domain specific region with a very precise working, i.e., has only a single use. However, more recently evidence has suggested that the working of this area is employed for multiple uses.

Chapter 3 – A Neural Reuse Perspective of the FFA: The Search for a Working

Broadmann's map is comprised of 52 regions and has played a fundamental role in giving form to cognitive function for more than one hundred years. From a modularity perspective, Broadmann's map fits relatively well, providing well-articulated structural boundaries in which functions can be conveniently placed. However, as cognitive neuroscience has progressed, I would argue that this map has become less useful and actually presents a hindrance to advancements across a wide range of disciplines whose interests lie in understanding cognitive function.

The problem is that, while useful in sectioning the brain into structurally distinct regions, Broadmann areas can tempt researchers to assign them function. For example, this could occur when researchers observed brain activity in a Broadmann area for a

given task, (e.g., fMRI data reveal activity in Brodmann area 44 during language production tasks, therefore the function of Brodmann area 44 is language production) or, a Brodmann area can be examined from the outset of an experiment to seek a new function that might be associated with it or to confirm a function that has been proposed through past research; in this case, research into other uses of the region. In both scenarios it is assumed a certain brain region is uni-functional, and that function is likely confined to the boundaries of a relatively small brain area, e.g., an anatomical module, however, opinions on this vary as modules are considered by some to span large distances across the brain.

This is not to say that Brodmann's map is in some way erroneous, rather it has reached its limit of utility, much like Newtonian physics required relativity and quantum mechanics to further advance our understanding of the universe in its smallest and largest capacities.

Likewise, PET and MRI technology have historically been used (according to their capabilities) to examine *macro-level* tasks, such as word identification, face and object recognition, semantic processing, memory, etc (Bechtel, 2003). As discussed earlier, Bechtel (2003) has flagged this level of examination for concern:

...the identification of these tasks with one brain area is partly an artifact of the sensitivity of the instruments and the statistical measures invoked in earlier studies, which only revealed areas with the greatest increase in activation, and the use of the subtractive method.

Therefore, we can't see what is going on beneath the surface – here, the surface being the limits of current neuroimaging technology. It should of course be underscored that next generation PET/MRI fusion technology has already begun its descent into finer

grained regions of the brain. It is conceivable that as this technology evolves further, more can be learned about cognitive workings, in terms of their activity and range of uses.

However, in the absence of this precise, next-generation technology, I have developed a three-pronged approach to explore the working of a particular brain region. First, uses associated with a single modality (e.g., vision – recognition of faces vs objects, for example) can be examined in an attempt to infer a working from a particular brain region. I will use this approach to demonstrate that the working of the fusiform face area (FFA) is not face-specific (single domain) by seeking out non-face uses (multi-domain), e.g., animals or objects (Within Modality Analysis). Second, multi-modality uses can be explored for regions that seem to play a strong role with respect to one modality in particular. In this regard, I will examine the literature for possible links between the FFA, which appears to have a strong observed use in face recognition, and other modalities, audition in particular (Between Modalities Analysis). Third, broad pools of neuro-imaging data can be scanned to identify behavioural and/or cognitive tasks for which the FFA has shown strong to moderate activity (Between Domains Analysis). These three analyses will be carried out in sections 14, 15 and 16, respectively.

Prior to exploring these three methods in an attempt to identify the working of the FFA, I will spend some time reviewing the literature associated with this brain region. In particular, I will focus on research that leads to and has reinforced the view that the FFA is domain-specific for human faces, and other recent research that suggests that the FFA plays a general ‘expertise’ role that reaches beyond faces.

Section 13: Fusiform Face Area

In the previous Chapter, significant attention was dedicated to assessing whether workings are fixed and at what grain size we should focus on when exploring a given working. For my examination of the FFA, I will assume that the working of the FFA is fixed and that the FFA is the appropriate grain size to be looking at in order to learn something new about cognitive function, from a neural reuse perspective. In Chapter 2, homology was also identified as something that can be useful for narrowing in on a working of a particular brain region. This topic will not be covered at this time but should be the focus of future research relating to the reuse of an FFA working. As described above, the focus of this Chapter is to explore the type of information processed by the FFA (domain specific versus domain general, and across modalities) and the connectivity of the FFA to other workings that, together, support a broad range of uses. However, first I will provide background on what is known about the general functioning of the FFA, with respect to its role in face recognition.

A portion of the fusiform gyrus, called the FFA, has been shown to respond relatively strongly to faces (Rhodes et al, 2004; Rossion, 2008); both visually observed and imagined faces, albeit with a reduced level of activity for the latter (O'Craven & Kanwisher, 2000). The FFA is located on the lateral side of the fusiform gyrus of the temporal lobe and typically is larger in the right hemisphere.

While the evidence strongly suggests that the FFA plays a role in face recognition, it is not clear what exactly that role is. There are at least two possible ways that we can describe the FFA's involvement in face recognition. First, the neuroanatomical boundaries of face recognition could lie entirely within the FFA, i.e., the

FFA is sufficient for face recognition. If this is the case, the FFA, or a suite of components within the FFA, gives rise to face recognition. In this scenario, input could flow from several other brain areas but if connections that innervate the FFA were severed, face recognition would not occur. Another scenario could be imagined as follows: the FFA plays a role in face recognition, i.e., the FFA is necessary but not sufficient for face recognition.

If the FFA were sufficient for face recognition we would expect that 1) the FFA is the end point of a chain of lower level inputs; and 2) FFA activity is associated with a representation of a face, not face parts. If the FFA were necessary but not sufficient for face recognition, we would expect that the FFA processes information related to an attribute of a face, e.g., a face part or the location of face parts, or semantic information associated with a face, e.g., athlete, relative, and race. The current state of evidence on this issue is mixed, however it does help direct our attention for future research and provide important information regarding how ‘face recognition’ should be defined.

Zhang et al (2012) have expanded on the work of Rhodes et al (2004) and others in an attempt to determine whether representation of faces in the FFA is holistic or parts-based. In other words, Zhang et al (2012) have attempted to determine if featural information (e.g., eyes, nose, mouth) of faces and configural information (spatial relation of facial parts) of faces are processed together (holistic) in the FFA, or if they are processed independently (parts-based). Their research revealed that the absence of first-order face configurations (i.e., the vertical location of parts) weakened the ability to discriminate between face parts (Zhang et al, 2012). The discrimination of face parts was heightened for two serially presented faces (one shortly after the other) when the first-

order configuration was present (Zhang et al, 2012). This suggests that the FFA represents faces in a holistic way, i.e., information about facial features and facial configuration is processed together, rather than independently. These findings suggest that the FFA might be necessary and sufficient for face perception – the end of the line for face processing, where faces are represented.

Based on findings from this same study, Zhang et al (2012) also conclude that the occipital face area (OFA), of the inferior occipital cortex, is associated with the parts-based representation of faces. This conclusion was based on results that showed activation in the OFA regardless of configural differences (Zhang et al, 2012). This complements findings that have shown that transcranial magnetic stimulation (TMS) of the OFA interferes with the ability to discriminate face parts, but not configurations (Zhang et al, 2012). The results of a number of other studies suggest that early visual areas bypass the OFA and project directly to the FFA (Rossion, 2007). Together, these findings strengthen the case for the FFA being the end location at which visual inputs come together to produce face recognition and therefore suggest that it is sufficient for face recognition.

However, the story might be more complex than the findings of Zhang et al. (2012) lead us to believe. Ishai (2011) describes a more complex cognitive architecture for face recognition than that of Zhang et al, and puts forward a distributed cortical network model, which describes the FFA as an essential piece of a broader architecture, i.e., realization occurs at multiple sites, including the FFA. According to Ishai (2011), this cortical network is comprised of the inferior occipital gyrus (IOG) and lateral fusiform gyrus (FG); the superior temporal sulcus; the amygdala and insula; and the

inferior frontal gyrus. Each area is thought to add its own unique role to our perception of a particular face, e.g., identification of an individual involves processing in the IOG and FG, whereas the inferior frontal gyrus (IFG) is thought to process semantic information, amygdala for facial expression, the superior temporal sulcus (STS) for gaze and speech-related movement, and orbitofrontal cortex (OFC) for aesthetics (Ishai, 2011).

Ishai (2011) proposes that the cortical network for face perception has “core” regions and “extended” regions, where core regions process invariant facial features and extended regions process changeable features of the face. The core system is said to be hierarchical, wherein the IOG influences both the FG and STS. The FG has been shown to strongly influence the extended system, including the amygdala, IFG and OFC (Ishai, 2011). Interestingly, the STS of the core system does not seem to influence the extended system (Ishai, 2011). Further evidence suggests that the FG plays the role as mediator within the overall system. Indeed, increased coupling has been seen between the FG and amygdala when individuals view emotional faces (Ishai, 2011). Likewise, coupling increased between the FG and OFC when famous faces were viewed (Ishai, 2011). Based on her model, Ishai (2011) also predicts strengthened coupling between other regions from hypothesized tasks or inputs, e.g., viewing animated faces would increase connectivity between the STS and IFG/OFC.

Ishai’s hypothesis that a system is required for face recognition, rather than a single region, is further supported by research conducted by Avidan et al (2005), which describes individuals with prosopagnosia who display normal levels of activity in the FFA. Further to this, Avidan & Behrmann (2009) demonstrated that congenital prosopagnosia results from decreased connectivity between core and extended systems.

Rossion (2007) is also generally supportive of Ishai's proposed system approach, however he extends the system further to include the anterior temporal cortex, citing research involving a number of prosopagnosic patients with damage to this area. Rossion (2007) asserts that the OFA, FFA and anterior IT (excluding the pSTS) are necessary components for face identity processing. Rossion (2007) notes that it is possible that the FFA's role is to categorize the visual input of a face but that this depends on the OFA to add the precision required to differentiate one face from another, i.e., to individualize the input.

In terms of how we might interpret the research above from a neural reuse perspective, it is tempting to imagine the FFA as a working that performs within-category disambiguation. If this were the case, it would not be surprising that it has been labeled as a face recognition module, as information about faces is perhaps the most salient type of visual information (in an evolutionary context) and therefore the most deserving of within-category disambiguation and the most likely to dominate the activity of this brain region. In the following section, I will describe why the FFA has historically been viewed as a domain specific module, and will then explore other research which points to a multi-domain specificity description of the FFA wherein its working supports a range of 'expert' uses; a small step from the 'category disambiguation' proposed earlier.

Section 14: The Role of the FFA: Competing hypotheses (A Within-Modality Analysis – Single Domain versus Multi-Domain)

The Single Domain View

In this section, I will examine whether the FFA is specialized for one specific type of input (e.g., human faces) or multiple types of input (e.g., human faces, animal faces,

animals, objects) within the visual modality. The FFA has largely been thought of as a domain specific brain region; a module evolved to perceive faces. The earliest claims that the FFA might be critical for processing faces are derived from observations of neuropsychology. Lesions to the FFA region have long been understood to produce cognitive deficits specific to face recognition, a condition that has been named prosopagnosia. Sergent et al (1992) were the first to link the FFA to prosopagnosia by means of PET examination of face-recognition and object-identification tasks. In addition to identifying the FFA as critical for face-recognition, they also provided evidence for the dissociation of face and object processing; the latter associated with activity in the left occipito-temporal cortex (Sergent et al, 1992). Notably, lesions to the right ventral occipito-temporal cortex have been shown to be sufficient to cause prosopagnosia (Ishai, 2011).

There are two types of prosopagnosia: acquired and developmental. The former is produced by damage to relevant brain regions, while in the latter form the individual never acquires the ability to recognize faces. In terms of acquired prosopagnosia, Gainotti & Marra (2011) describe two distinct types of face identification impairment. Damage to the occipital face area and the FFA are associated with the most common forms of prosopagnosia (i.e., inability to recognize a face by its assembly of parts), whereas face impairment deficits in those with temporo-occipital damage are thought to result from an inability to access semantic information associated with the visual stimuli (Gainotti & Marra, 2011).

Kanwisher et al. (1997) were among the first to describe the FFA as a domain specific region specialized for the perception of faces. Kanwisher et al (1997) used a

number of tests to demonstrate that the FFA responds significantly more to faces than other objects, including images of houses and hands.

More recently, Rhodes et al. (2004) examined three competing hypotheses about the function of the FFA: face specificity, which proposes that FFA processes faces only; individuation, where the FFA is activated to differentiate within a homogenous class; and expert-individuation, wherein the FFA is activated only for differentiation of a homogenous class with which we are highly familiar. The fMRI activation findings of this study support the face-specificity hypothesis: the FFA showed greater activation to faces than to other stimuli; and processing for Lepidoptera (a large class, including moths and butterflies) occurred in a different region of the fusiform gyrus from that of face processing, which occurred in the FFA (Rhodes et al, 2004).

While the above evidence does suggest a domain-specific function of the FFA, it is conceivable that the results might actually be indicative of a working that has developed a bias toward the recognition of faces due to their high importance and ample inputs. Indeed, the FFA working could appear to be domain-specific because faces have preempted its activity, but might in fact be domain general under the right circumstances, e.g., circumstances where a certain type of salient input is presented often. The following section provides evidence to support this hypothesis.

A Multi-Domain View

An alternative to the FFA being domain specific for face recognition is that it is domain general for recognition of anything visual. But why has it come to be thought that the FFA is specialized to recognize faces? Two possible explanations come to mind:

- 1) As alluded to above, perhaps faces outcompete other visual inputs in the use for the

FFA's working; and 2) maybe researchers have failed to look for non-face evidence because the modularity perspective has caused them to be content with the interesting link between this area and face recognition.

The idea that faces might outcompete other visual stimuli for use of the FFA working is in line with findings reported by Anderson (2010), which suggest that when multiple uses of a working are required at the same time, one will likely be muted or negated. Further to this, if a use of a given working becomes specialized early in development, it could become harder to put the working to another use in the future (however, it should be noted that the opposite could also occur – see Section 10 discussion of Anderson & Penner-Wilger, 2012). Thus, it could be that the saliency of faces, combined with our constant exposure to them, has dominated activity, and perhaps predisposed, the FFA for this particular use. If this is the case, however, it should be possible to train the FFA to become more sensitive to other visual inputs, with increased exposure. This view clearly presents a challenge to those such as Kanwisher (2010) who contend that the neural circuitry of the FFA is not influenced by experience, but rather is entirely, or almost entirely, genetically prescribed to respond to faces.

With respect to the second idea, that researchers have become distracted with a face-recognition module, there is little that can be done to verify this claim. One option would be to search the vast fMRI data that report FFA activity and check which types of tasks were presented in these studies. Studies that used tasks that are obviously consistent with face-recognition would clearly suggest that the goal of the researchers was to explore this particular use. However, if the search revealed tasks that were unrelated to face-recognition, this would suggest that the researchers were either

interested in examining other uses of the FFA (and were successful) or they were interested in a particular cognitive function that happened to use the FFA. The former scenario would shed light on the extent to which the modularity has biased our understanding of the FFA, while the latter would provide new insights into the working of the FFA. Given the importance of this latter scenario for my work, I will address this in more detail in Section 16.

For the remainder of this section, I will provide a summary of the most significant findings in support of a domain general working of the FFA. This research suggests that the FFA's working is one that precisely categorizes highly detailed visual input, i.e., that it is a visual expertise working. Indeed, this 'expertise hypothesis' provides interesting insights on how the FFA might be understood, by outlining a much broader range of capabilities than that is centered on differentiation rather than simply recognition.

The Expertise Hypothesis

Bukach et al (2006) present findings that support a domain general explanation of the neural machinery responsible for face recognition. In particular, they highlight the plasticity of these brain regions and the influence of experience in molding them.

Bukach et al (2006) found that both the FFA and OFA became more active in perceiving Greebles (novel object that have general similarities and specific, yet subtle, differences) following ten hours of Greeble training. Bukach et al (2006) further determined that the increased activity of the FFA was associated with holistic processing and not configural processing. These findings are consistent with Ishai's cortical network model of face perception and also with the finding of Rossion (2007) cited above.

Complementary findings demonstrate that performance of this domain-general holistic processing is proportional to the amount of training, and that the expertise might be acquired in a particular spatial direction that expands with training duration. Indeed, Bukach et al (2006) found that longer Greeble training periods were associated with holistic processing of the entire Greeble, whereas shorter training periods resulted in processing of Greeble segments only, e.g., the upper half. Similarly, other work by Bukach describes that some prosopagnosics process some parts of the face and not others. It is conceivable that these prosopagnosics suffered damage to parts of the OFA, given that this area has been linked to parts-based representation of faces. However, it is difficult to extend these findings to plasticity within the domain general holistic processing network because we do not understand how training effects the components of this network.

Research examining competition among different domains further supports the hypothesis that so-called face perception circuitry is domain general and its neural resources are co-opted to perform multiple functions. For example, N170 responses for faces are reduced when Greeble experts are asked to simultaneously process Greebles (Rossion et al, 2004). Moreover, car experts exhibit reduced N170 values for faces when they process cars at the same time (Gauthier et al, 2003). As expected, the level of reduction is proportional to the level of expertise, where greater expertise for car discrimination is associated with greater N170 reductions for face processing (Gauthier et al, 2003).

Section 15: Similar Domain General Uses Across Sensory Modalities (A Between Modalities Analysis)

As evidence indicates, human expertise for face-recognition, along with potentially many other visual inputs, employs a working whose operation is carried out within the FFA. It is conceivable that this highly specialized, yet domain general, working zone might have analogous structures in other brain locations that support similar functions for other modalities. Moreover, it is also possible that the FFA's working is employed by other modalities to support different uses. This hypothesis has been presented as far back as 2001, when Pascual-Leone and Hamilton (2001) put forward the idea that although the brain appears to be functionally organized by modality, this is actually a result of functional suppression that occurs over development, which is consistent with both Johnson's IS framework and Anderson's IDS framework.

Much as visual cognitive processing attends to shapes, colours and position, auditory processing uses frequency, pitch and amplitude to attribute categorical aspects such as gender, age, race and emotion. Further to this, it can be argued that for those with a full set of senses, audition is the second most valued of the senses, next to vision, from both an evolutionary and developmental perspective. Thus, a search for analogs might be best directed at the audition modality.

A relatively small amount of cognitive, psychological and neuroscience research has been undertaken in this area; however, the findings provide interesting insights that can inform our understanding of neural reuse generally, and expert holistic processing specifically. Chartrand et al. (2008) equate voice recognition with the "auditory face", where voices rapidly represent unique individuals much as faces do. Evidence for this is observed in individuals who develop an expertise for auditory recognition beyond our

ability to recognize familiar voices (Chartrand et al. 2008). For example, experienced musicians and bird watchers have been shown to exhibit a profound ability to identify a particular type of instrument (among many) or bird, respectively (Chartrand et al. 2008). Interestingly, some musicians with expertise in recognizing music also exhibit advantages in processing pitch and contour in language (Chartrand et al. 2008), which suggests that an expertise working might be being reused across higher-level functions.

There is evidence that the neuroanatomical sites associated with voice recognition are located in the same general area as those responsible for face-recognition. The superior temporal sulcus has been described as the core area for voice recognition; however, the anterior temporal pole, fusiform gyrus and posterior cingulate gyrus (pCG) have also been linked to this function (Arnott et al. 2008). As in face-recognition, the right hemisphere plays a more prominent role than the left (Arnott et al. 2008).

Much like prosopagnosics, individuals who have damage to anterior portions of their temporal lobe have difficulty recognizing people from their voice (Arnott et al. 2008). Interestingly, the pCG and FFA have been shown to play functional roles in both modalities. Indeed, the pCG is activated during familiar face recognition and familiar voice recognition (Arnott et al. 2008). Further to this, activity in this area occurs in response to voices that are personally familiar to the individual, but not in response to famous (e.g., celebrities') voices (Arnott et al. 2008). Based on these findings, it has been proposed that the pCG combines input from multiple domain-specific areas in order to represent familiarity (Arnott et al. 2008).

As discussed in previous sections, there is substantial scientific evidence implicating the FFA as an important neuroanatomical site for face recognition. Activity

in the FFA has also been demonstrated to occur during voice recognition (Arnott et al. 2008). Interestingly, Von Kriegstein et al (2005) found that activation of the FFA during voice recognition occurs only when the voice is familiar to the individual. Further to this, functional connectivity analysis has revealed that the FFA connects directly with the superior temporal sulcus, the core site for voice recognition (Arnott et al. 2008). In addition, face recognition is strengthened when an individual was previously exposed to the voice that is associated with the face (von Kriegstein et al. 2005). Taken together, these findings not only suggest that face and voice recognition interact at the behavioural level, but that the FFA might provide a multimodal working for person recognition (i.e., the ability to recognize individuals not only by their face, but other aspects as well, such as their voice), or perhaps more generally ‘individuation expertise’, and which is developed through experience and learning (i.e., the ability to differentiate between individuals within a particular category – humans, animals, plants or objects, for example). While this working is clearly linked to person recognition, this should not prevent us from looking for an even more domain-general working, such as recognition, or expertise, of salient things.

It is conceivable that FFA activity in voice recognition could be the result of visually imagining an individual’s face as the result of a verbal prompt, however fMRI data reveal that FFA activity during voice recognition is primarily influenced by projections from the superior temporal sulcus, the core region for voice recognition (von Kriegstein et al. 2005). This early binding process is inconsistent with what would be expected from a top-down activation of the visual system, triggered by a voice (von Kriegstein, 2005).

Moreover, FFA activity associated with voice recognition has also been observed in developmental prosopagnosics who have never been capable of perceiving or imagining faces (von Kriegstein et al. 2005). However, research conducted by Arnott et al (2008) provides somewhat contradictory findings as they report that an individual with extensive damage to the FFA (among other regions) maintained some level of voice recognition (Arnott et al. 2008). In this individual, activity associated with voice recognition was observed in the left posterior temporal regions, dorsal posterior cingulate and medial parietal regions (Arnott et al. 2008).

This introduces an important consideration. While I have focused exclusively on the reuse of one particular configuration of neurons to support multiple uses, neural reuse can also be examined in at least one other way. A particular working could possibly be located in multiple brain regions where it supports the same or different uses. This could include workings at multiple active sites, i.e., a working that is regularly performing the same or different use at multiple sites within the brain, or workings that are not active but can become active (or become established) if a brain region that houses a particular working is damaged. For this latter scenario, a neighbouring brain region, or even the alternate hemisphere might replace a damaged working. Smith & Sugar (1975) present interesting findings in which a young girl who underwent a left-hemispherectomy went on to develop superior language and intellectual abilities as an adult. Thus, taking this into consideration, the findings of Arnott et al. (2008) might simply be explained by the plasticity of the brain, and therefore might not contradict the findings of von Kriegstein et al. (2005).

Section 16: Between-Domains Analysis and Connectivity of the FFA

As we strengthen our understanding of how the mind works and as neuroimaging technology advances, we will be both better positioned to examine the brain at a finer grained level than now, and to identify and understand the interaction of low-level operations. A great deal has already been learned about the forward, backward and lateral projections within the visual system (see Koch, 2004). Bechtel (2003) describes the map of the visual system as the most comprehensive map of a complex system that we have. Its projections have been examined from its efferents at the eyeball up through each cortical layer. While this system relates to a single sensory modality, similar endeavors have been and will be targeted at higher-level cognitive processes and those that are multimodal in nature.

An abundance of research is being conducted with the objective of mapping the brain's billions of connections. The NIH Connectome project is one large-scale initiative that supports the development of connectivity maps of the brain (NIH). Connectivity maps can provide a good starting point from which possible reuse of cognitive workings can be inferred. Indeed, understanding how a particular region interacts with other regions and which functions these other regions are associated with can provide insights into the working of a particular area.

In this section, I will use this approach in an attempt to further clarify the working of the FFA. I will examine moderate to high connectivity of the FFA with other brain regions and will identify the functions that have previously been linked to these regions, with the intention of inferring a common working.

While connectivity of the FFA was touched upon in Sections 14 and 15, in the context of face and voice recognition, this section provides a more comprehensive and detailed description, across a broad range of domains. The findings will be used to hypothesize about the actual evolved working of the FFA. In other words, I will explore beyond the high level function typically associated with this area, i.e., face recognition, in an attempt to identify a working that evolved in a context much different than the one we are familiar with today, and has been coopted to support new challenges over our evolutionary history. Just as Broca's area is thought by some to be a hypersequential processor (working) that has been put to use to support the function of language, among other things, this exploration will attempt to uncover the working in the FFA that supports the face recognition use.

The left fusiform gyrus connects strongly with the left bank of the superior temporal sulcus, left entorhinal cortex, left and right inferior temporal cortex, left parahippocampal cortex and right fusiform gyrus (NIH). It also connects weakly with the left lingual gyrus, left middle temporal cortex, and left frontal pole (NIH). Table 1 below provides a functional overview for each connection site so that they can be more easily considered against one another.

Table 1: FFA Connectivity to other brain regions

Brain Region	Associated Functions	Comments
<i>Strong Connectivity</i>		
Superior temporal sulcus	Joint attention – perception of where others are looking; and perception of biological motion.	In individuals without autism, active when hearing human voices.
(Left) Entorhinal cortex	Memory (spatial memory in particular) and	First region to be affected by Alzheimer's disease.

	navigation.	
Left/Right inferior temporal cortex	Visual object recognition	Impairment in visual recognition or agnosia, including faces, colours, animals, objects.
(Left) Parahippocampal cortex	Visuo-spatial processing, episodic memory, encoding and recognition of environmental scenes.	Active when viewing landscapes, cityscapes, rooms. Damage to this area can result in a scene recognition deficit, leaving object/face recognition in tact.
<i>Weak Connectivity</i>		
(Left) Lingual gyrus	Vision processing (especially letters). Activates during memorizing and maintaining images of human faces. Also active in selective visual attention tasks.	Thought to play a role in analysis of logical conditions (i.e., logical order of events). Lesions cause visual memory impairments and are associated with visual snow syndrome.
(Left) Middle temporal gyrus	Associated with contemplating distances, face recognition and accessing word meaning while reading.	Lesions may result in alexia or agraphia.
(Left) frontal pole	Memory recall and executive function.	High connectivity with multi-modal regions.

Table 1 demonstrates that the FFA is connected to areas that play a role in very diverse uses. There are a number of insights that one might draw from this information. First, strong connectivity is narrowly associated with visual processing (i.e., recognition and attention) and visual memory (i.e., spatial and episodic). This is consistent with what has been asserted in early sections: the FFA is involved with recognition/differentiation of visual input. Inferring from the evidence provided in previous sections, it appears that a *within-category disambiguation processor* is the common working of the FFA. Thus,

the involvement of memory here is not surprising given that memory would be critical for ensuring the long-term aspects of recognition. Indeed, recognition could not be recognition without a past experience to refer to.

What is most interesting about the strong connectivity regions is that the superior temporal sulcus has been shown to be inactive in autistics in response to human voice input. Autism is also thought to be associated with visual processing deficits, including an inability to maintain eye contact. Taking these two pieces of information together, it would be interesting to examine the FFA in autistics in order to determine a number of things: First, given that visual deficits are associated with autism, we might expect that FFA activity would be altered due to a decrease in visual input. This could result in a diminished capacity for face recognition in autistics; evidence does in fact exist to support this (Klin et al, 1999). Whereas face recognition deficits in prosopagnosics generally occurs from damage to the FFA and/or its surrounding regions, this deficit in autistics likely arises due to an information block that prevents the FFA from performing this particular use.

If autistics were not using their FFA for face recognition, it would be especially interesting to study its level of activity when other tasks are being performed. It is conceivable that because the FFA is not preoccupied with faces, it could be more available for other tasks of *within-category disambiguation*. Indeed, this might explain why some autistic savants possess a heightened ability to disambiguate between days of the week on which events occurred many years ago, or attribute precise statistics to sports figures or teams. In other words, the FFA in autistics could be occupied with non-visual information disambiguation, rather than face processing.

With respect to weak connectivity with the FFA, regions that also become active during memory tasks (likely working memory, given its support of executive function) and visual recognition are seen, as was observed in the strong connectivity regions. What stands out is the role the middle temporal gyrus and lingual gyrus play in reading and processing letters, respectively. Informationally, letters, words and phrases can be thought of as similar to faces. Much like faces, words represent something that is determined by parts that are shared and similar. Indeed, words share letters as faces share eyes, ears, noses and mouths. Moreover, the letters that words are comprised of are similar in that they are lines bent in different configurations, and like noses, eyes, ears and mouths differ in their configuration. It is conceivable that the FFA working combines with lingual gyrus and/or middle temporal gyrus to disambiguate words based on the letters they include, or disambiguate sentences based on the words they include.

Conclusion and Next Steps

In this project I have attempted to demonstrate four things: 1) modularity has provided a useful perspective by which to understand how the mind works; 2) modularity has limits in terms of how thoroughly it can explain cognitive function; 3) neural reuse provides a new perspective on cognitive architecture and has more explanatory value than modularity perspectives; and 4) when applied to the FFA, the neural reuse perspective exposes operational elements that are masked, or simply not considered, when modular approaches are employed to explain this area.

Given how our understanding of the mind has developed over the past century and half, it is not difficult to see how modularity established its place above other

perspectives. It is in our nature to deconstruct things, be it a biological system, a machine, or a brain, into their component parts and assign a role to each. We are most comfortable doing this at the highest level, as it becomes mechanically or conceptually difficult to do this at lower levels, e.g., when attempting to explain the function of a cell, it is easier to describe it in terms of its organelles, rather than the sub-functions within each organelle. This might be part of the reason why DNA was discovered more than two hundred years following the discovery of the cell nucleus. The same can be said about the deconstruction of the brain; it takes time to move from an explanation of a hemisphere to that of a lobe to that of a region, etc.

Further to this, neuropsychology has established that lesions to certain regions of the brain are consistently associated with the loss of specific functions (e.g., damage to Broca's area severely limits one's ability to produce language). While this evidence has been useful in helping us to assign structure to function, it has also primed our thinking about cognitive architecture toward a modular picture. Indeed, recognizing that a function is lost when a specific region of the brain is damaged fits well with our understanding of things as built from parts that each play a unique role and fit together like pieces of a puzzle. Once this relationship is established, future experiments focused at the given brain region (structure) tend to be framed around the particular use (function). For example, once Broca's area was established as a "language area" future experiments largely explored language-type tasks. It would seem that this provided a barrier to exploring other uses that Broca's area might be involved in, or responsible for.

Understanding the brain's architecture as one that involves workings that can be reused to support multiple uses provides greater explanatory value (i.e., its ability to fit

the facts, generate new predictions, and generate interesting questions) than modular perspectives. Indeed, if Broca's area is in fact a hierarchical sequencer, this tells us more about this region of the brain than a modular view that restricts its use to that of language production. One explanation is clearly more thorough than the other.

Although the neural reuse perspective is compelling in its potential to better explain how the brain is organized and how it supports function, it is relatively new and hasn't yet been tested sufficiently to become a theory that is widely subscribed to amongst cognitive scientists. Further studies will be required to determine whether this perspective can be put into practice. The work that has been done in identifying possible uses of Broca's area has given some credibility to the neural reuse perspective, however more evidence of this type is needed in order to support the proposition that neural reuse is an architecture that is consistent throughout the brain, the cortex in particular. This is why I have endeavored to examine the FFA through a neural reuse lens.

Based on my examination of the FFA from a neural reuse perspective, it can be said that this region of the brain is likely not domain specific, but rather operation specific, i.e., it is comprised of a definable working that supports multiple uses. It appears that this working supports multiple uses within the visual modality (e.g., recognition of faces and objects), and there is also some evidence to suggest that it plays a similar role in the auditory modality. Together, these findings suggest that rather than a domain specific face recognition module, the FFA might actually be better described as a *within-category disambiguation processor*. Assuming this to be true, we could imagine that the FFA has evolved to differentiate between similar yet different information that is important for our evolutionary fitness, thereby discriminating one grouping of input from

another grouping of similar input based on differences in relatively small detail. Further research should be carried out to explore any involvement of the FFA in auditory tasks, particularly voice, or more general sound, recognition to determine if this proposed working of the FFA can cross modalities.

Clearly there is an evolutionary advantage to being able to differentiate between similar things, be it faces, cars, or voices. Indeed, identifying the difference between a friend and an enemy could be the difference between life and death. Similarly, identifying the difference between a small, red poisonous berry and a berry that looks similar but provides nourishment could lead to a similar advantage. It is not necessarily important to identify the original selective pressure, or pressures, that led to the formation of this working; rather, it's the possibility that this working evolved and has subsequently been co-opted for multiple uses, each of which provide some sort of advantage (and possibly some disadvantages).

More research is also required to explain how the FFA becomes specialized for certain types of input. Anderson (2014) has proposed an intriguing hypothesis by way of his *Interactive Differentiation and Search* framework. While Anderson has put forward a plausible explanation for how workings are formed, in terms of the operation they carry out, and which uses they support, it has not yet been subject to the scientific and philosophical scrutiny that all claims of this magnitude must undergo. Further study into how neurons form relationships with one another and to what extent these relationships can change with time and with different inputs will go a long way in verifying or refuting Anderson's claims.

Based on knowledge of brain areas that are connected to the FFA, it seems that the FFA is linked to a broad range of cognitive capacities, including memory, attention, logic processing, reading, recognition and executive function. Although the FFA is connected to these brain regions, we know very little about when the FFA is active with these areas. It is important to note here that while simultaneous activation of the FFA and its connected areas is useful, what is more useful is that we know when these areas are supporting a common task. More research should be undertaken to identify tasks for which the FFA is activated in unison with each of its connected regions. The results from this type of study would provide a much stronger starting point than the one that I have used in my analysis, which depended simply on connectivity rather than simultaneous or sequential activation. In addition, it would target relevant tasks from which inferences about the working of the FFA can be made.

I hope that what I have found has in some way contributed to the discussion. I consider this topic to be one that could have significant impact across many areas of cognitive science. Likewise, unpacking neural reuse further and exploring its limits could greatly benefit researchers in neuroscience, computer science, linguistics, and philosophy. Each of these domains are well positioned to add value to the current theory, and could play a meaningful role in the scrutiny that the neural reuse perspective will need to undergo if it is to establish its place in cognitive science.

References

- Anderson, M. (2010) Neural reuse: A fundamental organizational principle of the brain. *Behavioral and Brain Sciences*, 33(4), 245-313.
- Anderson, M., & Penner-Wilger, M. (2013) Neural reuse in the evolution and development of the brain: Evidence for developmental homology? *Developmental Psychobiology*, 55(1), 42-51.
- Anderson, M. (2014) *After Phrenology: Neural reuse and the interactive brain*. MIT Press.
- Arnott, S.R., Heywood, C.A., Kentridge, R.W., & Goodale, M.A. Voice recognition and the posterior cingulate: An fMRI study of prosopagnosia. *Journal of Neuropsychology*, 2(1), 269-86.
- Avidan, G., Hasson, U., Malach, R., Behrmann, M., (2005) Detailed exploration of face-related processing in congenital prosopagnosia: 2. Functional neuroimaging findings. *Journal of Cognitive Neuroscience*, 17(7), 1150-1167.
- Avidan, G., Behrmann, M. (2009) Functional MRI reveals compromised neural integrity of the face processing network in congenital prosopagnosia. *Current Biology*, 19, 1146-1150.
- Barrett, H.C. (2005). Enzymatic computation and cognitive modularity. *Mind and Language*, 20(3), 259-87.
- Barrett, H.C., & Kurzban, R. (2006) Modularity in cognition: Framing the debate. *Psychological Review*, 113(3), 628-47.
- Bates, E. (2001) Modularity, domain specificity and the development of language. *In Philosophy and Neuroscience* (Bechtel et al), 134-51.
- Bechtel, W. (2003) Modules, brain parts, and evolutionary psychology. In S. J. Scher and F. Rauscher (eds.) *Evolutionary psychology: Alternative approaches*. Dordrecht: Kluwer.
- Bechtel, W, (2008) Mechanisms in cognitive psychology: What are the operations? *Philosophy of Science*, 75, 995-1007.
- Bennett, C. M.; Miller, M. B.; Wolford, G. L. (2009) "Neural correlates of interspecies perspective taking in the post-mortem Atlantic Salmon: An argument for multiple comparisons correction. *NeuroImage*, 47, S125.
- Bergeron, V. (2007) Anatomical and functional modularity in cognitive science: shifting the focus. *Philosophical Psychology*, 20(2), 175-95.

- Bergeron, V. (2008) *Cognitive architecture and the brain: Beyond domain-specific functional specification*. Unpublished doctoral dissertation, University of British Columbia, Vancouver, British Columbia, Canada.
- Bergeron, V. (2010) Neural reuse and cognitive homology. *Behavioral and Brain Sciences*, 33(4), 268-9.
- Bergeron, V. (2015) Functional independence and cognitive architecture. *British Journal for the Philosophy of Science*, doi: 10.1093/bjps/axv005.
- Blackmore, D. (2013) Why it is wrong and what it can teach us anyway. Unpublished masters thesis, Carleton University, Ottawa, Ontario, Canada.
- Bukach, C.M., Gauthier, I., Tarr, M.J. (2006) Beyond faces and modularity: The power of an expertise framework. *TRENDS in Cognitive Science*, 10(4), 159-66.
- Carruthers, P. (2006) *The Architecture of the mind: massive modularity and the flexibility of thought*. Oxford University Press.
- Chartrand, J.P., Peretz, I., & Belin, P. (2008) Auditory recognition expertise and domain specificity. *Brain Research*, 1220, 191-98.
- Chiao, J.Y., Harada, T., Oby, E.R., Li, Z., Parrish, T. & Bridge, D.J. (2009) Neural representations of social status hierarchy in human inferior prefrontal cortex. *Neuropsychology*, 47(2), 354-63.
- Clune, J., Mouret, J.B., and Lipson, H. (2012) *The evolutionary origins of modularity*. Proceedings of the Royal Society. B. 280: 20122863.
- Confer, J.C., Easton, J.A., Fleischman, D.S., Goetz, C.D., Lewis, D.M.D., Perilloux, C., & Buss, D.M. (2010) Evolutionary psychology: Controversies, questions, prospects, and limitations. *American Psychologist*, 65(2), 110-126.
- Davies, M. (2010) Double dissociation: Understanding its role in cognitive neuropsychology. *Mind and Language*, 25(5), 500-40.
- Deacon, T. (1997) *The Symbolic Species: The co-evolution of language and the brain*. New York: W.W. Norton & Company.
- Dehaene, S. (2005) Evolution of human cortical circuits for reading and arithmetic: The “neuronal recycling” hypothesis. In: From monkey brain to human brain. Ed. S. Dehaene, J.R., Duhamel, J.R., Hauser, M.D., & Rizzolatti, G, 133-57.
- Dekker, T.M., & Karmiloff-Smith, A. (2010) The importance of ontogenetic change in typical and atypical development. *Behavioral and Brain Sciences*, 33(4), 271-2.

Deroy, O. (2010) *Modularity and perception*. in M. Matthen (ed.) Oxford Handbook of Philosophy of Perception, Oxford University Press.

Donnarumma, F., Prevede, R., Trautteur, G. How and over what timescales does neural reuse actually occur? *Behavioral and Brain Sciences*, 33(4), 272-3.

Fodor, J. (1983) *The modularity of mind*. Cambridge, MA: MIT Press.

Gainotti, G., & Marra, C. (2011) Differential contribution of right and left temporo-occipital and anterior temporal lesions to face recognition disorders. *Frontiers in Human Neuroscience*, 5, 55.

Gallese, V. (2008) Mirror neurons and the social nature of language: The neural exploitation hypothesis. *Social Neuroscience*. 3(3-4), 317-33.

Gauthier, I., Curran, T., Curby, K.M., & Collins, D. (2003) Perceptual interference supports a non-modular account of face processing. *Nature Neuroscience*, 6, 428-32.

Glenberg, A.M., & Kaschak, M.P. (2002) Grounding language in action. *Psychonomic Bulletin & Review*, 9(3), 558-65.

Hirschfeld, L.A., & Gelman, S.A. (1994) *Mapping the mind: Domain specificity in cognition and culture* New York, NY: Cambridge University Press.

Hurley, S.L. (2008) The shared circuits model (SCM): How control, mirroring, and simulation can enable imitation, deliberation, and mindreading. *Behavioral and Brain Sciences*, 31(1), 1-58.

Immordino-Yang, M.H, Chiao, J. Y. & Fiske, A. P. (2010) Neural re-use in the social and emotional brain. *Behavioral and Brain Sciences*, 33(4), 275-276.

Ishai, A. (2011) A cortical network for face perception. In: *New Frontiers in Social Cognitive Neuroscience*. Editors: R. Kawashima, M. Sugiura, T. Tsukiura. 73-81. Tohoku University Press, Sendai.

Johnson, M.H. (2001) Functional brain development in humans. *Nature Reviews Neuroscience*, 2, 475-83.

Johnson, M.H. (2011) Interactive specialization: A domain general framework for human functional brain development? *Developmental Cognitive Neuroscience*, 1(1), 7-21.

Junge, J.A., & Dennett, D.C. (2010) Multiuse and constraints from original use. *Behavioral and Brain Sciences*, 33(4), 277-8.

- Kanwisher N, McDermott J, Chun MM. (1997) The fusiform face area: a module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience*, 17(11), 4302–11.
- Kanwisher, N. (2010) Functional specificity in the human brain: A window into the functional architecture of the mind. *Proceedings of the National Academy of Sciences of the United States of America*, 107(25), 11163-70.
- Karmiloff-Smith, A. (1992) *Beyond modularity: A developmental perspective on cognitive science*. MIT Press.
- Katz, P.S. (2010) Comparative studies provide evidence for neural reuse. *Behavioral and Brain Sciences*, 33(4), 278-9.
- Klin, A., Sparrow, S.S., de Bildt, A., Cicchetti, D.V., Cohen, D.J., & Volkmar, F.R. (1999) A normal study of face recognition in autism and related disorders. *Journal of Autism and Developmental Disorders*, 29(6), 499-508.
- Lieberman, D., Tooby, J., & Cosmides, L. (2003) Does morality have a biological basis? An empirical test of the factors governing moral sentiments relating to incest. *Proceedings of the Royal Society, B*, 270, 819-26.
- Luria, A.R. (1973) *The working brain: An introduction to neuropsychology*. Penguin Books Ltd.
- Machery, E. (2007) Massive modularity and brain evolution. *Philosophy of Science*, 74(5), 825-38.
- Marr, D., & Poggio, T. 1976. *From Understanding Computation to Understanding Neural Circuitry*. Artificial Intelligence Laboratory. A.I. Memo. MIT.
- NIH. www.humanconnectomeproject.org
- Nieder, A., & Miller, E. K. (2004). In S. Dehaene, J. R. Duhamel, M. Hauser & G. Rizzolatti (Eds.), *From monkey brain to human brain*. Cambridge, Massachusetts: MIT Press.
- O’Craven, K.M., & Kanwisher, N. (2000) Mental imagery of faces and places activates corresponding stimulus-specific brain regions. *Journal of Cognitive Neuroscience*, 12(6), 1013-23.
- Over, D.E. (2003) *Evolution and the psychology of thinking*. Psychology Press, Taylor and Francis Group.
- Pascual-Leone, A., & Hamilton, R. (2001) The metamodal organization of the brain. *Progress in Brain Research*, 134, 427-45.

- Patel, A.D. (2003) Language, music, syntax and the brain. *Nature Neuroscience*, 6, 674-681.
- Paterson, S., Brown, J., Gsodl, M., Johnson, M., & Karmiloff-Smith, A. (1999) Cognitive modularity and genetic disorders. *Science*, 268, 2355-8.
- Piazza, M., Izard, V., Pinel, P., LeBihan, D., & Dehaene, S. (2004) Tuning curves for approximate numerosity in the human parietal cortex. *Neuron*, 44(3), 547-555.
- Poldrack, R.A. (2006) Can cognitive processes be inferred from neuroimaging data? *Trends in Cognitive Sciences*, 10(2), 59-63.
- Prinz, J.J. (2006) Is the mind really modular? In Stainton, Robert J. (Ed), *Contemporary debates in cognitive science*, 22-26.
- Ritchie, J.B., & Carruthers, P. (2010) Massive modularity is consistent with most forms of neural reuse. *Behavioral and Brain Sciences*, 33(4), 289-90.
- Rhodes, G., Byatt, G., Michie, P.T., & Puce, A. (2004) Is the fusiform face area specialized for faces, individuation, or expert individuation? *Journal of Cognitive Neuroscience*, 16(2), 189-203.
- Rossion, B., Kung, C.C., & Tarr, M.J. (2004) Visual expertise with nonface objects leads to competition with the early perceptual processing of faces in the human occipitotemporal cortex. *Proceedings of the National Academies of Science*, 101(40), 14521-14526.
- Rossion, B. (2008) Constraining the cortical face network by neuroimaging studies of acquired prosopagnosia. *Neuroimage*, 40, 423-26.
- Sergent, J., Ohta, S., & MacDonald, B. (1992) Functional Neuroanatomy of face and object processing. *Brain*, 115(1), 15-36.
- Schubotz, R.I., & Fiebach, C.J. (2006) Integrative models of Broca's area and the ventral premotor cortex. *Cortex*, 42, 461-463.
- Smith, A & Sugar, O. (1975) Development of above normal language and intelligence 21 years after left hemispherectomy. *Neurology*, 25(9), 813-8.
- Sperber, D. (2001) In defense of massive modularity. In Dupoux, E. *Language, Brain and Cognitive Development: Essays in Honor of Jacques Mehler*. Cambridge, Mass. MIT Press, 47-57.
- Sperber, D. (1994) The modularity of thought and the epidemiology of representations, In L. Hirschfeld & Gelman (Eds.), *Mapping the Mind*, 39-67. Cambridge University Press.

- Thorpe, S., Fize, D., & Marlot, C. (1996) Speed of processing in the human visual system. *Nature*. 381(6582), 520-2.
- Tooby, J. & Cosmides, L. (1992) The psychological foundations of culture. In J.H. Barkow, L. Cosmides, & J. Tooby (Eds.), *The Adapted Mind*, 19-136. Oxford University Press.
- Tovee, M. J. (1994) How fast is the speed of thought? *Current Biology*, 4(12), 1125-7.
- Von Kriegstein, K., Kleinschmidt, A., & Giraud, A.L. (2005) Voice recognition and cross-modal responses to familiar speakers' voices in prosopagnosia. *Cerebral Cortex*, 16, 1324-22.
- Zhang, J., Li, X., Song, Y., & Liu, J. (2012) The fusiform face area is engaged in holistic, not parts-based, representation of faces. *PLoS ONE* 7(7): e40390. doi:10.1371/journal.pone.0040390.