Synthese (2007) 159:347–371
DOI 10.1007/s11229-007-9234-1

Why neuroscience matters to cognitive neuropsychology

Victoria McGeer

Published online: 26 September 2007

Abstract The broad issue in this paper is the relationship between cognitive psychology and neuroscience. That issue arises particularly sharply for cognitive neurospsychology, some of whose practitioners claim a methodological autonomy for their discipline. They hold that behavioural data from neuropsychological impairments are sufficient to justify assumptions about the underlying modular structure of human cognitive architecture, as well as to make inferences about its various components. But this claim to methodological autonomy can be challenged on both philosophical and empirical grounds. A priori considerations about (cognitive) multiple realisability challenge the thesis on philosophical grounds, and neuroscientific findings from developmental disorders substantiate that challenge empirically. The conclusion is that behavioural evidence alone is inadequate for scientific progress since appearances of modularity can be thoroughly deceptive, obscuring both the dynamic processes of neural development and the endstate network architecture of real cognitive systems.

Keywords Cognitive neuropsychology · Modularity · Multiple realisability · Neuroconstructivism · Williams syndrome (WS) · Autism (ASD) · Face processing

1 Introduction

This paper is focussed on the general issue of the connection between neuroscience and cognitive psychology. But it approaches this issue in the context of one particular subdiscipline of cognitive psychology: that of cognitive neuropsychology (CNP). One reason for focussing on this specialised subdiscipline is that cognitive neuropsychologists

V. McGeer (✉)
University Center for Human Values, Princeton University, 5 Ivy lane, Princeton, NJ 08544, USA
e-mail: vmcgeer@princeton.edu

have been very explicit about their methodology, and self-conscious about the arguments they make for and against the relevance of neuroscientific work to their enterprise. Thus, it is an appropriate site for teasing apart the philosophical and empirical considerations that go into taking a particular stand on this general theoretical issue.

CNP emerged in its contemporary form in the wake of the cognitive revolution. As described by its practitioners, its overall aim is shared with that of cognitive psychology more generally—namely, to develop theories of the functional architecture of (normal) cognition on the basis of task-specific behavioural studies. What makes this subdiscipline distinctive, therefore, is not its aims, or even necessarily its methods; rather, it is distinctive by virtue of the kind of data to which it appeals—viz., atypical patterns of performance found in a variety of affected populations. In short, CNP aims to study normal cognitive (including motor and perceptual) processes by investigating how such processes can be impaired, either as a result of acquired brain damage or as a result of a developmental disorder (Caramazza and Coltheart 2006, pp. 3–4).

This characterisation of CNP presupposes what many of its practitioners actively affirm: that there is a clear disciplinary boundary between, on the one hand, cognitive psychology—the study of the functional organisation of (human) information-processing systems—and, on the other, cognitive neuroscience—the study of the neural systems subserving cognition. To characterise this distinction somewhat tendentiously, one discipline studies the mind, the other studies the brain (Coltheart 2002a). Of course, for many scientists, this disciplinary boundary is largely conceptual, since they engage in both activities. But despite such 'interdisciplinarity', there is continuing disagreement about how these two activities ought to relate to one another—a theoretical dispute that has deep implications for how the study of cognition and its disorders proceeds.

My objective in this paper is to consider the theoretical and empirical grounds for maintaining what has come to be known as "ultra"-cognitive neuropsychology (ultra-CNP): the view that the study of impaired behavioural performance, whether this comes about as a result of brain damage or a developmental disorder, provides sufficient evidence for making real and substantial progress on delineating the underlying architecture of normal cognitive processes. In particular, what ultra-CNP denies is that neuroscientific observations play any necessary role in this enterprise, potentially derailing the conclusions reached on the basis of behavioural studies. This paper argues that this fundamental tenet of ultra-CNP is deeply misguided—and it is misguided on both theoretical and empirical grounds. The argument unfolds in three sections. In Sect. 2, I explore the basic assumptions of ultra-CNP in more detail, arguing that because these basic assumptions have substantial empirical implications, proponents of ultra-CNP seem to face a fundamental challenge from those who insist that neuroscientific work must play a critical role in developing theories of normal cognition. In Sect. 3, I consider what a priori considerations proponents of ultra-CNP have raised to counter such a fundamental challenge—unsuccessfully in my view. And in Sect. 4, I show that empirical work on developmental disorders—in particular neuroscientific work—provides just the sort of evidence theorists need to question the basic assumptions of ultra-CNP.

2 Ultra-CNP and the implementation challenge

A fundamental tenet of ultra-CNP is that the mind is modular. Many versions of the modularity thesis have been debated in psychology and philosophy,[1] but proponents of ultra-CNP endorse the following basic picture: the mind is as an information-processing system, composed of a number of task-specific subsystems (modules)—e.g. for reading, for face recognition, for mental state attribution. Each of these subsystems is composed in turn of more basic subsystems (modules within modules). And these in turn may decompose into more basic subsystems still (modules within modules within modules), and so on—not ad infinitum, of course, but to a point where there are no more basic (cognitive) tasks into which component modules can fractionate (Block 1995; Coltheart 2002a). Importantly, many theorists agree that some modular subcomponents might be shared across different subsystems, in which case there would be partial overlap within a modular functional architecture. Still, what makes the architecture modular is the fact that many subcomponents of such modules will not be shared—they will be uniquely dedicated to the (higher-order) task they serve (face recognition or whatever).

This picture is important for the kind of inferences that can be drawn from neuropsychological data once combined with three other key assumptions. The first, standardly called the "universality" assumption, is that human brains are qualitatively similar in respect of their cognitive architectures. The second key assumption, called "fractionation" or "subtractivity", is that modules function "relatively autonomously", allowing them to be selectively impaired either through brain damage or faulty development. Finally, there is the "transparency" assumption, that when damage or alteration does occur to some modular structure, no significant alterations occur in other parts of the system so as to create what are, in essence, new processing structures (Caramazza 1986, 1992; Harley 2004a).

The overall effect of these assumptions is to license the methodology of ultra-CNP, which stands on two central principles. The first is that cognitive neuropsychologists should focus on the symptoms of individual patients, rather than group studies. The rationale for this approach is simple. If the mind is modular roughly in the way sketched above, then cognitive systems will fractionate into complex nested structures of modular components. Any combination of these modular components might conceivably be affected through lesioning or atypical development, leading to distinctive patterns of cognitive abilities and disabilities. By comparing a multiplicity of individual differences, it should be possible to build up a composite picture of underlying modular structure. For instance, after brain damage if individual A is able to do task α to a normal degree of competence but not task β (a single dissociation), then this provides some evidence that task β is at least partially accomplished by some modular

[1] In a contemporary context, the *locus classicus* is Jerry Fodor's *Modularity of Mind* (Fodor 1983). However, the picture of modularity presented here differs from Fodor's in some significant respects. For instance, under Fodor's conception, modules do not fractionate into submodules, since they are "not assembled". However, the aim in this paragraph is not to recapitulate Fodor's views, but to present a picture that is widely shared in cognitive neuropsychology (for discussion, see Coltheart 2002a, pp. 140–143). In this context, I will also use the term 'module' and 'subsystem' interchangeably, though many theorists have used these terms in differentiating ways.

component not required by α. Of course other explanations for this dissociation are possible. Perhaps task α requires less working memory than task β, or is less susceptible to some other processing constraint, making it easier for the system to do. But suppose there is also another individual B who is able to do task β but not α, then the evidence from A's and B's performances put together (the famed double dissociation) now constitutes much stronger evidence that tasks α and β are accomplished by separate modules. Other explanations are still possible. But, given the constraints of the background assumptions, such explanations are significantly more complicated so less persuasive than the separate module hypothesis. Hence, by combining data from a number of different individuals on a range of cognitive tasks, the hope is to build up a catalogue of dissociations that will lead in turn to an ever more refined (composite) picture of the functional architecture of the various cognitive systems constituting a (normal) human mind.

The second principle on which the methodology of ultra-CNP stands is implicit in the example discussed above. Behavioural studies are the primary source of data for inferences to modular structure: i.e., patterns of dissociations, and in particular double dissociations, found amongst a variety of individuals on a range of cognitive tasks. If two theories propose different modular hypotheses as a way of making sense of extant data, then the primary—some would say, only—means of determining whether either or both these theories is false is to discover some disconfirming pattern of behaviour. As Caramazza and Coltheart remark, "... a single patient can refute some hypothesis about cognitive architecture by yielding a pattern of data that according to that hypothesis could never occur ..." (Caramazza and Coltheart 2006, p. 7).

Of course, this may be overstating the case somewhat. Even friends of ultra-CNP acknowledge that recalcitrant data is more likely to produce modification than refutation: auxillary hypotheses can always be added to save any particular theory from sinking under a single hit (Harley 2004b; Lakatos 1974). But more substantial concerns are in the offing. Van Orden and colleagues have argued that using behavioural data to identify double dissociations is inevitably an interpretive enterprise, compromised by two levels of circularity (Van Orden et al. 2001; see too, Bishop 1997). First, a reading of the data will depend on particular candidate theories of modular structure. That is to say, from the perspective of theory A, particular patterns of behaviour will look like cases of pure dissociation, supporting a particular modular hypothesis; but they will not appear so from the perspective of theory B. For instance, the phenomenon of 'deep dyslexia' is regarded by some as a pure syndrome implicating the existence of a (partly intact) lexical module; by others, it is regarded as a mixed syndrome with different modular implications. Without some way out of this circularity, it is hard to see how behavioural data alone can adjudicate among these different modular possibilities. Second, any such reading of the data will depend on endorsing the general picture of modularity described above, including the assumptions of universality, fractionation and transparency. In this case, it is hard to see how behavioural data can support *any* inferences to modular structure unless these assumptions are independently justified.

Proponents of ultra-CNP have defended the approach on two related grounds. In the first place, they say, the circularities so identified are no worse than the sort of circularities that plague any scientific enterprise, since they simply reflect a general problem with the theory-ladenness of data:

> In the context of cognitive neuropsychology this implies that descriptions of deficits are not independent of the theories that guide the observations we deem to be relevant. That is, the features of performance that will be analysed will be precisely those our theoretical perspectives guide us to seek, the measures of performance we will use will be determined by the theoretical contrasts of interest, and the general description of deficits will be determined by the theories we are willing to contemplate. However the relevance of this observation in the context of the evaluation of the validity of inferences in cognitive neuropsychology is mysterious. Granted that data are theory laden – a feature of observations in all empirical sciences – we are still led to ask whether the implications of this fact are particularly problematic for cognitive neuropsychology (Caramazza 1992, p. 87).

If this point is correct—that the circularities discussed above are just a particular instance of the more general problem of the theory-ladenness of data—then the second ground of defence seems fitting: It is that the proof of any scientific enterprise is in the pudding—that is, in the case of CNP, we should assess whether the methods used have produced any genuine progress in understanding the actual processes of human cognition. Correspondingly, we should not be too bothered by *possible* defects of this approach if no concrete case can be made for any actual defect—i.e. where the approach can be shown to retard rather than enhance real empirical progress (Caramazza 1992, p. 93).

Naturally this raises the question of what would count as real empirical progress in understanding the actual processes comprising human cognition; what would count, in other words, as an adequate cognitive theory delineating the functional components of the human cognitive system. Kossylyn and van Kleeck have suggested that such a theory must satisfy "two strong requirements": (1) the theory must be "computational adequate": it must provide an explicit specification of the component subsystems (or modules) along with a "prescription for how these subsystems work together to produce the overall capacity"; and (2) it must have a "structural realization": it requires ". . . a mapping of the putative subsystems and their interconnections onto the brain" (Kosslyn and Van Kleeck 1990, pp. 391–392). In Kosslyn and van Kleeck's view, the second requirement of specifying how the brain putatively accomplishes a given functional task is particularly critical—indeed, they suggest it is "what distinguishes the cognitive neuropsychology approach from that of conventional cognitive psychology" (p. 392). Whether or not they are correct in this suggestion, their fundamental concern seems well taken:

> Computational adequacy alone is of little interest to cognitive neuropsychology without evidence that the processes are actually realised in neural hardware. The goal is to provide a theory of the functional organization of the brain, not simply a concise or useful description of behaviour.

Call this the implementation challenge. It suggests that ultra-CNP, so far as it advocates developing cognitive theory without bothering about the implementation details, is importantly misguided. At best, it may just be strategically misguided—misguided in consequence of downplaying the second requirement as a matter of practical priority.

But, at worst, it may be theoretically misguided—misguided in consequence of understanding the second requirement to place no deep constraints on developing cognitive theories; that cognitive theorizing can, and indeed must, proceed independently of whatever is discovered at the neural level.

If these charges are right, then the kind of defence offered by proponents of ultra-CNP against complaints of circularity and in favour of their method of generating functional theories, cannot stand. In the first place, critics will point out that real empirical progress can only be measured in terms of how well a functional theory meets both computational adequacy *and* implementation requirements. And, in the second place, the kind of circularities mentioned above flow from the need to make certain assumptions about design constraints on the functional architecture of the mind/brain. Naturally, behavioural data alone will not be sufficient to test these design constraints, since it is interpreted in terms of them. But to claim that all data is theory-laden in this way is to discount the possibility that the design constraints themselves are open to review through a better understanding of neural systems. This is a strong position, and we need good reason to accept it against those who argue to the contrary. In the following Sect. 3, I consider what a priori grounds there might be for maintaining such a position. Finding these to be inadequate, I will then turn in Sect. 4 to a consideration of how these assumptions may be challenged on the basis of neuroscientific evidence.

3 Theoretical issues—ultra-CNP and its a priori bid for the methodological autonomy of CNP

Proponents of ultra-CNP are committed to the claim that behavioural observations ". . . stand on their own in the development of a meaningful cognitive science"—that is to say, ". . . developments in cognitive science concerning the computational structure of cognitive processes can proceed independently of neuroanatomical observations" (Caramazza 1992, p. 85). This is a bid for the methodological autonomy of cognitive psychology from neuroscience. There are, however, weaker and stronger versions of this basic position, and both of them merit consideration. The more extreme version—an ultra-ultra-CNP—ups the ante on methodological autonomy. The claim is not just that cognitive psychology *can* proceed independently of neuroscientific work; it is that cognitive psychology has *no choice* but to proceed independently of such work: Since there is nothing in principle that neuroscientific studies could contribute to cognitive theory, it must be developed on the basis of behavioural studies alone. By contrast, a more moderate ultra-cognitivist position—I will call it barely-ultra CNP—concedes that neuroscientific work might, and probably will, be helpful in the development of cognitive theory; nevertheless, such work contributes nothing that is essential. In principle, if not in practice, meaningful cognitive theory can be developed on the basis of behavioural studies alone.

To focus the target of this discussion more precisely, it's worth reviewing the points on which the more extreme and more moderate versions of ultra-CNP concur. First, they both agree that the path to a mature cognitive science is conceptually a two-stage process, which involves different disciplinary activities. Stage one, the proper province of cognitive neuropsychologists, involves the development of an adequate functional

theory of human cognitive systems, one that specifies "the processing structure of hypothesized components and the types of representations they compute" (Caramazza 1992, p. 87). Once this stage is complete, cognitive neuroscientists will then be in a position to investigate how the functional components specified in such a theory are realised in neural structure; this is stage two. As Coltheart says, "in order to localise the modules of a cognitive system, one must first know what the systems modules actually are. So we must begin with a model and then seek to do localisation research" (Coltheart 2004, p. 23).

Clearly there is an implicit expectation contained in this program: that a satisfying implementation account will be forthcoming once an adequate cognitive theory has been developed. But what if it is not? Stone and Davies have argued that any theorist who is committed to a "moderate neuron doctrine" will accept that neuroscientific observations must here impose a limiting constraint on cognitive theory in the following sense: "the theory would have to be rejected if there were no neurobiological story consistent with it" (Stone and Davies 1999, p. 850). Are proponents of both versions of ultra-CNP committed to the moderate neuron doctrine?

I am not so sure. Certainly it would be odd for cognitive neuropsychologists to continue to endorse a cognitive theory that turns out at stage two to have is no clear or consistent realisation account—at least so long as they are committed realists who regard cognitive processes as instantiated *somehow* in the physical brain. However, as Stone and Davies caution, ". . . defensible claims about the theoretical and practical priority of the cognitive level sometimes tip over into something more extreme and implausible; namely, the claim that neurobiology is strictly irrelevant to cognitive psychological theorising" (Stone and Davies 1999). Proponents of ultra-ultra-CNP may be teetering on this edge. After all, if they insist that cognitive theoretical work not only *may* abstract from debates at the neuroscientific level (as per barely-ultra-CNP), but further *must* abstract from these debates, then it is hard to see how neuroscientific evidence can suddenly provide a testing constraint at the end of the day. Or to put this concern another way, if neuroscientific evidence is held to provide an ultimate test for cognitive theory, what stands in the way of its also providing "ongoing checks, balances and inspirations" in the on-going development of cognitive theory (Churchland and Churchland 1996)? Mere consistency seems to suggest that these two possibilities go together. Hence, one of the costs of embracing ultra-ultra-CNP may be to give up on the moderate neuron doctrine. This need not entail a rejection of physicalism; but it suggests a more limited, instrumental role for cognitive theory in the sense derided by Kosslyn and van Kleeck—namely, of providing "a concise or useful description of behaviour". I presume this is not a consequence they would be happy to embrace. Still, I will set this issue aside to focus more directly on the explicit methodological claims.

I come not to the main topic of this section: ultra-CNP's official position on whether and to what extent neuroscientific work has a role to play in the development of cognitive theory itself—i.e., in stage one, as this is conceptualised above. This is the point of doctrine on which ultra-ultra-CNP and barely-ultra-CNP explicitly diverge. I will begin by considering arguments in favour of the stronger view, concluding that those on offer are very far from persuasive. This suggests that the more moderate view really constitutes an appropriate fallback position. But my conclusion is that it does

not, because it fails to acknowledge the ways in which neuroscientific evidence must play a critical shaping role in developing cognitive theory.

The clearest statement of ultra-ultra-CNP comes from Max Coltheart (but see also, Mehler et al. 1984):

> . . . [F]acts about the brain do not constrain the possible nature of mental-information processing systems. No amount of knowledge about the hardware of a computer will tell you anything serious about the nature of the software that computer runs. In the same way, no facts about the activity of the brain could be used to confirm or refute some information-processing model of cognition. This is why the ultra-cognitive-neuropsychologist's answer to the question, 'Should there be any 'neuro' in cognitive neuropsychology?' is 'Certainly not; what would be the point? (Coltheart 2004, p. 22).

Although the computer analogy is suggestive, more needs to be said to justify such a strong position. There seem to be two rationales offered in its favour, the first serving as a kind of background consideration and the second given explicitly as a motivating set of reasons. I discuss each of these in turn.

The background consideration is linked to a central tenet of philosophical functionalism: that the functional architecture of information-processing systems is *supposed* to be specified at a level that abstracts away from physical details. This honours the fact that a given functional architecture *could* be realised in a multiplicity of physical ways (the multiple realisability thesis—or MR). And, indeed, this may explain why Coltheart and others are so keen to emphasise that embracing modularity at the cognitive level is not to embrace it at the anatomical level; and, in general, that there are no strict implications from how things are organised at the cognitive level to how they are organised at the anatomical level, or vice versa (Coltheart 2002a, 2004; Coltheart and Langdon 1998; see too, Mehler et al. 1984). The strict separation between these two levels of analysis may suggest that the sort of evidence one would need to adduce for a particular cognitive architecture has nothing to do with the sort of evidence one would need to adduce for a particular anatomical architecture; and, further, that the only kind of evidence that gives any real traction on cognitive architecture must be gleaned from the careful study of patterns of behaviour, especially as these may be selectively disrupted through a variety of impairments.

This point is familiar, but it is also precisely the target of Kosslyn and van Kleecks's critique. The purpose of developing theories that specify the functional components of cognition is not to give an analysis of how various tasks *might* be accomplished by the brain; it is to specify how such tasks *are* accomplished by the brain, so why not develop theories that are "consistent with the observed neuroanatomy and neurophysiology" from the get-go (Kosslyn and Van Kleeck 1990, p. 400)? Furthermore, Kossyln and van Kleeck worry that behavioural data can only go so far in winnowing down candidate theories. And this is surely true, both in practice and in principle. In practice it is true, because there will always be a limited amount of behavioural data available from which inferences must be made. And in principle it is true, because no amount of behavioural data could ever solve the problem. The reason for this is multiple realisability (MR): Only in this context, it is not MR at the level of physical

implementation that is of concern, i.e. the many possible ways of realising a particular algorithm in particular physical stuff. Rather it is MR at the cognitive level, i.e. the many possible ways of algorithmically computing a high-level cognitive task. (For a nice discussion of these two levels of MR, see Gerrans 2003b; and also Dennett 1987: "Instrumentalism Reconsidered".)

This last point deserves special emphasis and clarification. There has been so much attention paid to the phenomenon of physical-MR (hence to the problem of constraining cognitive theory from below) that the phenomenon of cognitive-MR tends to be overlooked, with the result that behavioural data comes to be treated as a kind of Archimedean point effectively constraining cognitive theory from above. But there is no such Archimedian point, even when the behavioural data includes patterns of impairment. The reason is that trade-offs can always be made between specifying an underlying cognitive architecture and specifying how that architecture has been impaired to produce the relevant pattern of atypical behaviour. Thus, behavioural data will always leave a number of candidate theories in the running (recall here the circularity concerns of van Order and colleagues mentioned in Sect. 2). Worse, it may well be that behavioural data biases theorists towards cognitive architectures that look to be more plausible from the top-down, as it were, but which have relatively limited plausibility from the point of view of explaining how such architectures could be realised in neural structures (Dennett 1987, pp. 75–76). For these various reasons, it seems like misplaced philosophical purism to claim that cognitive theory should only be answerable to behavioural data. Neuroscientific work may not be definitive in the kind of constraints it places on cognitive theory, but nor is behavioural data itself. So it seems the appropriate lesson to learn from these observations is not that one source of data should be privileged over others, but that all sources of data should be used in a combinatorial way to provide multiple levels of constraint on cognitive theory (Kosslyn and Van Kleeck 1990; cf. Churchland and Churchland 1996).

A second and more explicit argument offered in defence of ultra-ultra-CNP trades on the idea that the functional level of analysis has a certain priority over the neural level in the order of explanation. One aspect of this priority idea is that any explanation of how the brain manages to accomplish various cognitive functions will necessarily presuppose an account of what the cognitive functions are that it accomplishes. The logic of this point is impeccable; but from here, proponents of ultra-ultra-CNP move rapidly to the idea that pursuing an account of cognitive functions—i.e. developing cognitive theory—is something that must be done first and/or independently of neuroscientific implementation work. Can such a move be justified?

As far as can be seen, proponents of ultra-ultra-CNP rest their conclusion on two supporting claims: (1) unless one has a fully worked out cognitive theory, "one will not know what to look for at the physical instantiation level"; and (2) the only way to get to a fully worked out cognitive theory is to focus exclusively on behavioural data (Coltheart and Langdon 1998; see too, Coltheart 2002a, 2004). But neither of these claims is defensible, and together they certainly do not support any strong position on the methodological autonomy of CNP.

The first claim rests on the assumption that the order of discovery must follow the order of explanation. But there is nothing in the explanatory priority thesis that requires such a view. While it may be perfectly correct to insist that a complete explanation of

how the brain implements various cognitive functions will require an account of what those functions are, it is another issue entirely as to how the nature of those functions will be discovered. So, barring any substantial considerations to the contrary, theorists may consistently endorse the explanatory priority thesis while yet insisting that the order of discovery must involve ". . . a jagged co-evolution of cognitive science and neuroscience, moved forward by multiple cross-adjustments at the level of results and theory" (from an anonymous reviewer, quoted approvingly in Caramazza 1992, p. 85).

The second claim, that the only way to develop cognitive theory is to focus exclusively on behavioural data, goes to the heart of the issue. Proponents of ultra-ultra-CNP seem to embrace this claim because of the following epistemological worry: that "it can be very hard to understand what a system is doing if one's only information about it is at the physical level" (Coltheart and Langdon 1998, p. 150). However, if this is the only rationale offered in defence of the claim, it is hardly adequate to the task. Of course it can be very hard to understand what a system is doing if one's *only* information about it is at the physical instantiation level (my emphasis). But this is not an argument *for* the methodological autonomy of CNP. At most, it is an argument *against* the methodological autonomy of (cognitive) neuroscience. To suggest otherwise is to assume that the following options for developing cognitive theory are exclusive: either to use information from the physical instantiation level or to use information from the behavioural level. But, of course, this is not the case. The obvious and sensible methodological alternative is to use information from both levels of analysis in mutually constraining ways.

Notice how these reflections mirror the earlier reflections inspired by taking on board the real possibility of cognitive-multiple realisability. There it was pointed that it can be very hard to understand what a system is doing (at the level of specifying what algorithms it is 'computing') if one's only information about it is at the behavioural level. But, in parallel fashion, this was not an argument *for* an exclusive focus on information at the physical instantiation level—and, hence, *for* the methodological autonomy of cognitive neuroscience. At most, it was an argument *against* an exclusive focus on information at the behavioural level—and, hence, *against* ultra-ultra-CNP.

These two sets of considerations thus point to the same conclusion: that there is no in principle rationale for insisting that the development of cognitive theory ought to be pursued first and/or independently of understanding how cognitive functions are realised in the brain. On the contrary: they indicate that CNP and cognitive neuroscience ought to be pursued in tandem, with observations at both levels contributing in complementary ways to the development of a mature cognitive science. As the Churchlands nicely summarise, "theories at different levels quite properly function as ongoing checks, balances, and inspirations for theories at adjacent levels, both up and down" (Churchland and Churchland 1996, p. 221).

This alternative vision of substantial, on-going inter-level theoretical development is deeply opposed to the methodological principles articulated by ultra-ultra-cognitive-neuropsychologists; and it suggests that, so far as those principles hold sway in the actual practice of CNP, there needs to be some adjustment in the way researchers approach their discipline. But how much adjustment is really required? After all, even those who are sympathetic to the view that "behavioural observations stand on their own in the development of meaningful cognitive science" might concede that

neuroscientific work could contribute significantly to the development of cognitive theory; hence, in practice, they might be more than happy to encourage such inter-level research. This seems to be the position of what I have called 'barely-ultra-CNP', represented most explicitly by Caramazza, who, on the one hand, argues vehemently against Kosslyn and van Kleeck's indictment of a purely behavioural approach in CNP, and, on the other hand, repeatedly insists that neuroscientific work is important to pursue (e.g. Caramazza 1992, p. 85). Given this more liberal practical stance, why should it matter if barely-ultra-cognitive psychologists maintain, perhaps merely as a point of principle, that "behavioural observations stand on their own in the development of meaningful cognitive science"?

The concern this position raises is more subtle, but it may have consequences for how cognitive psychologists think about and pursue their research agenda. The question now is not about whether neuroscientific observations have any role to play in the development of cognitive theory, but about what sort of role such observations could play. In particular, do they have any power to challenge the assumptions about cognitive architecture discussed in Sect. 2, and so play an essential shaping role in the development of cognitive theory? Barely-ultra-cognitive neuropsychologists are sceptical about this; and this scepticism appears to derive, as in the case of ultra-ultra-CNP, from a misapplication of the idea that functional characterisations have explanatory priority over neurobiological observations in accounting for patterns of behaviour.

Consider, for instance, the trajectory of Caramazza's defence of ultra-CNP against the criticisms of Kossylyn and van Kleek. Beginning from the priority idea that it is functional characterisations rather than neurobiological observations that must pay their way, so to speak, in terms of explaining patterns of behaviour, Caramazza justly remarks that, though anatomical information might ". . . serve a useful role in clueing us in to different functional disorders, . . . it would have to have behavioural consequences of some sort. Otherwise, we would not have any grounds for concluding that the functional deficits are different in the first place" (Caramazza 1992). Yet, from this true and relatively benign methodological maxim, he draws a far more loaded conclusion—namely, ". . . that, in the final analysis, the relevant information for drawing inferences *specifically* about cognitive processing can ***only*** be patients' patterns of performance" (Caramazza 1992, p. 92, my emphasis in bold). Why *only* patients' patterns of performance? As we have seen from the earlier reflections on *cognitive*-multiple realisability, behavioural studies alone are not sufficient for drawing solid inferences about actual cognitive processing, so why not turn to neuroscientific evidence as a way of constraining possible interpretations? Caramazza's answer appears to be that neuroscientific evidence cannot aid in this particular endeavour. As he says, ". . . for one to maintain that the different possibilities for how function is realised in the brain would affect the nature of inferences from impaired performance for cognitive theory, it would have to be shown that depending on the specific assumptions one makes regarding the brain/cognition mapping, different conclusions *about cognitive theory* would be reached from the *same pattern of impaired performance*. It is not obvious to me that such an effort would meet with success" (Caramazza 1992, p. 90—emphasis in the original).

The problem with this argument is that it trades on a subtle mistake. It is not that different conclusions about cognitive theory would have to be reached from the "same

pattern of impaired performance". This *would* flout the sensible methodological maxim expressed above: that different functional proposals ought to have different behavioural consequences (however mild these may be). Instead, the argument depends on denying an altogether more substantial empirical possibility—namely, that depending on the specific assumptions theorists make regarding brain/cognition mappings, different conclusions about cognitive theory would be reached from the ***same behavioural evidence***. What makes this a substantial empirical possibility, as opposed to a mere methodological misstep, is due to the fact, already emphasised in Sect. 2, that 'patterns of impaired performance' and 'behavioural evidence' are not one and the same thing. Patterns of impaired performance are *interpretations* of behavioural evidence—and, importantly, such interpretations may be labile under different theories of cognitive architecture.

So even the milder form of (barely)-ultra-CNP finally depends on establishing one of two claims: either (1) that assumptions about cognitive architecture do not substantially affect interpretations of behavioural evidence after all, or (2) that neuroscientific evidence has no bearing on which assumptions theorists are justified in making about underlying cognitive architecture. The first claim is presumably a non-starter. As discussed in Sect. 2, proponents of ultra-CNP readily admit that the interpretations they make of behavioural evidence depends on three substantial assumptions regarding the structure of normal human cognitive architecture (universality), how it can be impaired (fractionation), and how it responds to such impairment (transparency). So this means that ultra-CNP must stand or fall on the remaining claim: (2) that neuroscientific evidence has no bearing on whether or not these or other assumptions are the appropriate ones to make. In this section, we have seen that there are no a priori considerations that count in favour of such a claim. However, perhaps it could be replaced by something a bit more modest—namely, the empirical claim that, so far at any rate, there is no neuroscientific evidence that bears on whether or not these or other assumptions are the appropriate ones to make.[2] This is the claim I will be disputing in the next and concluding section.

4 Testing the assumptions of ultra-CNP—evidence from developmental disorders

The two versions of ultra-CNP we have examined agree on this basic methodological principle—that behavioural studies "stand on their own" in providing the kind of data theorists need in order to make inferences to the underlying modular structure of normal cognitive processes. Indeed, compared with behavioural studies of normal subjects (involving response times, error rates and the like), studying the sometimes quite specific cognitive effects of naturally occurring "biological manipulations" is thought to reveal more or less directly how nature—in the form of cognitive architecture—is

[2] Interestingly, Coltheart raises this milder possibility only to explicitly reject it in favour of the more hardline position discussed earlier in this section (Coltheart 2004).

carved at its joints.[3] From this point of view, it should not matter how such naturally occurring manipulations come about: whether, for instance, by traumatic brain injury, by disease or by abnormal development. These are importantly different "distal" causes to be sure. But the thought is that, since varying "distal" causes may eventuate in the same underlying cognitive impairment, they can be ignored in terms of exploring what that cognitive impairment actually is; they can be ignored in terms of identifying the particular damaged or malformed cognitive module that serves as "proximate" cause of the observed atypical behavioural pattern. In sum, it is proximate causes, not distal causes, that are germane to the project of CNP (Coltheart 2002a, pp. 156–159; see also, Coltheart 2002b; Coltheart and Jackson 1998).[4]

Of course, there will be scientists who are engaged in a different kind of project—namely, the project of investigating the distal causes of particular developmentally induced cognitive abnormalities. But even for them, according to proponents of ultra-CNP, the work of (ultra-) CNP must serve as the critical first step. This is because an appropriate investigation of distal causes (i.e. the various processes, and their causes, that have led to the current situation) must surely depend on first identifying what needs to be explained—i.e. the proximate cause of some behavioural abnormality or, in other words, how precisely the relevant cognitive system of a given individual has been impaired (Coltheart 2002a, p. 159). This is just another application of the priority thesis discussed in Sect. 3.

Again, the logic of this basic position may seem impeccable; but the last ten years has seen an intensifying theoretical debate about whether or not cognitive neuropsychologists are right to assume that distal causes can be ignored in the investigation of currently existing cognitive impairments. In particular, a new breed of 'neuroconstructivists' argue that even if the cognitive 'end state' for acquired and developmental disorders looks rather similar in many respects from a behavioural perspective, the

[3] It may seem surprising that anatomical damage should 'carve nature at its cognitive joints', as if there were some precision to these accidents of nature. But, of course, there is no precision except in the following sense: anatomical lesions can have a number of specific cognitive effects, despite being accidental. In fact, it is precisely this phenomenon that theorists have used as an argument in favour of cognitive modularity. As modularists might say, what better explains the fact that 'untargeted' anatomical lesions have specific cognitive effects than that a number of different modules have been taken out at one go? Moreover, as discussed in Sect. 2, the distinctness of the resulting cognitive profile is taken to show, not only which modules have been affected, but also which modules must be present and operating normally to produce undamaged (a.k.a. typical) patterns of behaviour in healthy individuals. Of course, these are also the sort of inferences that are being questioned in this paper.

[4] As originators of these terms, Coltheart and Jackson most often use the distal/proximal distinction in the way that is indicated here—i.e. primarily to mark a temporal causal order of events: "when a system's operation is defective, there is not just one cause, but a chain of causes. The cause that is closest to the defective behaviour is the proximal cause; the other (more remote) causes are distal causes" (Coltheart 2002b, p. 157; see too, Jackson and Coltheart 2001). However, Coltheart has also used the distal/proximal distinction, not to mark a temporal ordering of causes per se, but rather to mark different levels of description of the self-same cause. For instance, in reference to prosopagnosia, he writes: "The proximal cause of this disorder is abnormality of one or more of the components of the cognitive system used for recognizing faces: it is cognitive neuropsychology's job to propose theories about what those components might be... The distal cause of this disorder is damage to the mechanisms of the brain that are involved in face processing, and that is the province of cognitive neuroscience, not of cognitive neuropsychology" (Coltheart 2002a, p. 157). To avoid confusion, I will restrict my use of the terms 'distal' and 'proximal' to the temporal ordering of causes, rather than as a way of marking different levels of realisation.

entire trajectory of development has been different in these two cases. This makes it likely that the 'end state' itself will also be very different, unless there is good reason to believe that abnormal cognitive development proceeds much like normal cognitive development after all. But there is no good reason for believing this, apart from endorsing a very strong thesis of genetically predetermined modular development. And this is a thesis that neuroconstructivists dispute. More to the point, neuroconstructivists argue that when behavioural evidence is taken on its own, it provides *misleading* support for this general thesis; and it's only when such evidence is combined with evidence from other sources, particularly from the neurosciences, that an alternative picture begins to emerge.

In sum, neuroconstructivists explicitly reject the central methodological position of ultra-CNP. In their view, the inadequacy of this position is no more clearly demonstrated than in cases of atypical development, where pockets of apparent behavioural normalcy can mask deeper functional differences between atypical and typical populations (Bishop 1997; Karmiloff-Smith 1998; Oliver et al. 2000; Thomas and Karmiloff-Smith 2002). In this section, I discuss the merits of the neuroconstructivist challenge through the consideration of a particular representative case: the capacity for face-processing in individuals with Williams syndrome (WS), as compared with both typically developing individuals and individuals with autism spectrum disorder (ASD).

WS is a genetically based developmental disorder that has both physical and cognitive manifestations. The cognitive manifestations are prima facie characterisable in terms of an uneven profile of 'spared' and 'impaired' cognitive abilities. Among those cognitive abilities that seem to be relatively spared are some aspects of language and social cognition, whereas those that are seriously impaired include spatio-constructive skills, numerical cognition, problem solving and planning (Bellugi et al. 1999; Karmiloff-Smith 1998; Wang et al. 1995). Many theorists have taken this uneven profile of cognitive abilities and disabilities to suggest an uneven profile of modular development that is commensurate with a developmental version of the assumptions of universality, fractionation and transparency discussed in Sect. 2 (these are sometimes referred to collectively as an assumption of "residual normality", e.g. by Thomas and Karmiloff-Smith 2002, 2003). That is to say, certain cognitive modules (for aspects of language use, for face processing) develop normally (explaining behaviour that tests within the normal range), whereas other modules (for numerical cognition, for visuospatial cognition) are seriously impaired or fail to develop at all (explaining behaviour that tests well below the normal range).

Further support for this theoretical approach is garnered from the range of 'spared' and 'impaired' capacities that characteristically occur in other developmental disorders, including autism (or ASD) (Baron-Cohen 1989; Leslie 1991), Gilles de la Tourette syndrome (Baron-Cohen 1998), specific language impairment (SLI) (Gopnik 1997), and various types of developmental dyslexia and dyscalculia (Temple 1997). Taken together, these disorders are said to provide the strongest possible evidence for developmental modularity thanks to the occurrence of 'dissociations' and even 'double dissociations' in their collective phenotypic outcomes. For instance, WS and SLI are held to constitute one such 'developmental double dissociation', with WS individuals showing a spared capacity for syntactic processing against a background of

multiply impaired cognitive capacities, and SLI individuals showing a specific deficit of syntactic processing against a background of generally spared cognitive abilities (Pinker 1999; Smith 1999; but for critical discussion, see Gerrans 2003a; Karmiloff-Smith 1998). Similar claims have been made about developmental dyslexia—indeed, with theorists using apparent 'double dissociations' in the acquisition of different putative modules of the normal adult reading system to recommend a more refined (provisional) typology of developmental dyslexias: e.g., 'developmental phonological dyslexia' versus 'developmental surface dyslexia' (Coltheart et al. 1993; Temple 1997). Hence, the picture of normal development reinforced by such evidence is one according to which there is predetermined epigenesis of dedicated domain-specific mechanisms; and the corresponding picture of abnormal development is one according to which genetic mutation or other forms of biological insult (e.g. maternal virus in utero) can selectively impair aspects of this predetermined epigenetic program.

But are these behavioural appearances to be trusted? Take face processing in WS, which may be profitably compared to face processing both in the normal case and in the case of individuals with ASD. In typical development, infants are initially drawn to face-like visual stimuli, progressively coming to devote more discriminating attention to specifically human faces in contrast to other face-like stimuli as well as to other objects. Infants with WS share this preference for face-like stimuli, even showing a disproportionate interest in faces compared to that of normal controls (Mervis and Bertrand 1997; Bellugi et al. 2000; Laing et al. 2002). By contrast, this bias seems to be completely lacking in individuals with ASD. Indeed, from an early age, autistic individuals show looking-patterns that are rather the reverse of what is normal: they spend more time looking at objects than at people; when looking at people, they tend to avoid looking at the face; and when looking at the face, they tend to avoid looking at the eyes (Osterling and Dawson 1994; Baron-Cohen 1997; Swettenham et al. 1998; Klin 2000). Thus, it seemed reasonable to surmise that face processing could well be impaired in individuals with ASD, though 'spared' in individuals with WS (Bellugi et al. 1997; Udwin and Yule 1991); and initial behavioural studies seemed to bear this dissociation out. WS individuals were found to score in the normal range on many standardised face recognitions tasks (Bellugi et al. 1990; Wang et al. 1995), whereas ASD individuals have shown impairments in a number of areas including: matching faces across different poses and images (Davies et al. 1994; Klin et al. 1999), incidental face learning (Boucher and Lewis 1992), memory for recently presented faces (Ellis et al. 1994), recognition of recently presented faces (Boucher et al. 2000), and detection of emotional expressions (Joseph and Tanaka 2003; Adolphs et al. 2001).

Such behavioural abnormalities in individuals with ASD point to substantive cognitive atypicalities that are bound to have clear neurological manifestations. ERP and fMRI imaging studies do in fact indicate less right-hemisphere laterisation and specialisation for face processing tasks (Schultz et al. 2000; Dawson et al. 2002; Pierce et al. 2001). One area of particular interest, hypothesised to activate in all face processing tasks in normal individuals, is the middle part of the right fusiform gyrus—the so-called 'Fusiform Face Area (FFA)' (Kanwisher et al. 1997). In individuals with ASD, there have been consistent findings of hypoactivation in this area, possibly explaining many of their face-processing abnormalities. Still, this can only be part of the story. For instance, more recent studies indicate that when ASD subjects are

instructed to attend explicitly to the internal versus peripheral features of the face (especially including the eye region), activation levels in the FFA are in the normal range (Hadjikhani et al. 2004; Dalton et al. 2005). This suggests that abnormal face processing in ASD must also involve other neural areas mediating attention to, interest in, and affective response to faces; hence, a full explanation of autistic face-processing abnormalities will have to look to the larger network within which the FFA is embedded.

This conclusion is consistent with the growing consensus amongst cognitive neuroscientists that normal face-processing occurs in a distributed neural system, primarily lateralised on the right side, but involving multiple cortical and subcortical areas. Various models for this system have now been proposed (e.g., Haxby et al. 2000; Elgar and Campbell 2001). However, a common hypothesis is that this system is constituted by (at least) two distinct, yet interacting, processing streams that involve separate regions of the brain. These streams can be seen as 'specialising' on different aspects of face-processing tasks—e.g. in Elgar and Campbell's (provisional) model, the "lateral" stream (involving middle temporal gyrus, superior temporal sulcus and ventrolateral parts of the prefrontal lobe) is more directly implicated in recognition of familiar faces, whereas the "medial" stream (involving parahippocampal gyrus, hippocampus, amygdyla and ventral orbito-frontal regions) is more directly implicated in learning new faces and for imbuing faces with significance. However, as Elgar and Campbell further caution, such specialisation should not be understood to operate in an excessively modular way: "Although we draw a distinction between a medial and a lateral system, they are assumed to operate together. Different face processing and recognition tasks are likely to be differentially dependent on each of these systems and their interactions. For example, whereas learning new faces may rely heavily on the medial circuitry, the lateral system may be recruited when the task requires us to compare the new face with one we already know. Furthermore, recognition of a familiar face involves not only the retrieval of semantic information (name, occupation), but also feelings—especially feelings of familiarity. Indeed, it is common to have feelings of familiarity about a face before precise identification evidence becomes available" (Elgar and Campbell 2001, p. 707).

The considerable cross-talk that occurs between these two networks makes teasing apart functional contributions of various 'parts' of the system a Herculean, perhaps even a Pickwickian task. In addition, given the distributed nature of this system, it seems likely that the various neuronal structures implicated in it will not be solely recruited (or recruitable) to face-processing tasks. For instance, even the FFA, which is the best known candidate for a face dedicated area and which certainly seems to serve as some sort of critical gateway in normal face processing, may not be exclusively *for* processing faces after all, even in normally developing individuals. Imaging studies indicate that this area is involved more broadly in making the kind of fine visual discriminations that allows for the within-category identification and recall of particular individuals, whatever the category in question (e.g., violins, birds, cars or simply nonface-like visual patterns). This is a mark of 'expertise'. Of course, for most of us, such 'expert' discrimination is rare outside the domain of faces, explaining why the FFA should be highly, and even selectively, tuned to faces in most individuals. Still, if this expertise view is correct, then it seems more appropriate to call the FFA

something that captures its true functional flexibility—e.g. as some theorists suggest, the "flexible fusiform area" (Tarr and Gauthier 2000; Gauthier and Logothetis 2000; but cf. Grill-Spector et al. 2004 for critical discussion).

While many details of the neuroscientific account of face-processing in normal individuals remain open to investigation and debate, the emerging picture of a complex widely distributed network is not one that seems particularly congenial to the characterisation of a cognitive capacity that is selectively 'spared'—or even 'impaired'—in the context of a developmental disorder. For instance, while it is true that many developmental disorders have highly specific and characteristic phenotypic outcomes, the underlying neural abnormalities are often substantial and diffuse. This is certainly true of both ASD and WS. In WS, for instance, there are abnormalities at the anatomical level (e.g. adult WS brain is 80% of the normal volume; total cerebral grey matter is significantly reduced; dorsal hemispheres show malformation; limbic structures and frontal cortex are also significantly reduced in size; and there is abnormal layering, clustering, orientation and size of neurons); in addition, there are also abnormalities in brain chemistry that, combined with structural abnormalities, are assumed to affect the timing and efficiency of neural processing (Bellugi et al. 1999; Rae et al. 1998; Grice et al. 2001; Mills et al. 2000). Given the breadth and depth of such abnormalities across many cortical and subcortical areas, it seems unlikely that the complex neural circuitry involved in processing faces could be unaffected. ERP studies support this surmise, with findings of waveform abnormalities in adult WS subjects that indicate both deviance and delay. In particular, there is overall less right hemisphere laterisation and less specialisation for human faces in upright orientation, as would be expected on a normal developmental trajectory (Mills et al. 2000; Grice et al. 2001). Indeed, these findings are more in keeping with the neural profile of autistic subjects discussed above, who also show signs of absent right-hemisphere laterisation and specialisation on face-processing tasks.

In sum, the neuroscientific evidence presents a two-part puzzle for the standard modular view of cognitive development and cognitive competency. In the first place, if similarities in the behavioural competencies between normal and WS individuals on face-processing tasks is to be explained by appeal to the pre-determined epigenesis of a modular capacity for face-processing, then it seems surprising that there would be evidence of much neural deviation in the way that capacity is realised in normal compared with WS individuals, as seemingly is the case. And secondly, if the behavioural *abnormalities* in autistic individuals are to be explained in terms of the maldevelopment of this modular capacity, then it seems surprising that a neural signature of this maldevelopment (e.g. lack of right hemisphere laterisation and apparent specialisation for face-processing tasks) should also be present in individuals with WS, where this modular capacity is supposedly preserved. How is this puzzle to be resolved?

The beginning of a resolution comes from returning to the behavioural evidence, only this time with an eye to looking in a more fine-grained way at both typical and atypical developmental trajectories. It turns out that WS individuals show subtle signs, not only of abnormal neural developmental, but also of abnormal cognitive development, casting doubt on the idea of a 'spared' normal competency for processing faces.

Looking first at typically developing individuals, it is well known that behavioural markers of a mature cognitive capacity for face-processing include a greater reliance

on configural over local (featural) processing: i.e., faces are identified, not on the basis of particular individual features (nose, eyes, chin, etc.), but on how these features are spatially related to one another and to the image as a whole. When faces are inverted, these higher-order configural properties are particularly disrupted, making recognition more difficult than when faces are presented in an upright orientation. Hence, the normal developmental trajectory is marked by an increasingly pronounced 'face-inversion effect' (Chung and Thompson 1995), as individuals shift from a more holistic method of face recognition (template-matching) in early infancy (de Schonen et al. 1989; Morton and Johnson 1991) through to more featural processing in young toddlerhood (Carey and Diamond 1977) and finally into the more configural style of processing that characterises skilled recognition. At the neural level, this developmental trajectory is marked by the fact that brain processes are initially very similar for processing human faces (in upright or inverted orientation), monkey faces and even other objects (de Hahn 2001), yet become increasingly more specialised and localised for upright human faces to the right-hemisphere fusiform gyrus (Passarotti et al. 2003).

The neural developmental profile of individuals with WS is, as we have seen, markedly different in just this respect: there is an apparent absence of right-hemisphere lateralisation and specialisation. But what about their cognitive profile? In fact, more refined behavioural studies have now shown impairments on face-processing tasks that require specifically configural processing (Karmiloff-Smith 1997; Karmiloff-Smith et al. 2004). In addition, WS individuals are like individuals with ASD in showing a reduced face inversion effect (Hobson et al. 1988; Deruelle et al. 1999), with further studies indicating that a general tendency towards a more piecemeal (or featural) processing style is a characteristic of both atypical populations (Happé 1999; Frith and Happé 1994; Deruelle et al. 1999). Thus, with respect to face processing in particular, it appears that individuals with WS achieve their relatively good performance on face discrimination and face memory tasks by using a different cognitive strategy than that of normal individuals—and, indeed, a cognitive strategy that is perhaps representative of a more general inability, readily apparent in the visuo-spatial domain, to integrate features into a configural whole (Karmiloff-Smith et al. 2004).

While these more recent behavioural findings do seem to gibe more naturally with the neuroscientific data, arguing for atypical cognitive structures in both WS and ASD, this conclusion generates a further puzzle that must be addressed. If it's true that WS individuals are more like ASD individuals in the way they process faces or indeed other sorts of visual stimuli, then why is it the case that ASD individuals do so poorly on many face processing tasks in comparison with WS individuals? Further, why do WS individuals do so poorly on many visuospatial tasks in comparison with ASD individuals, who on these tasks even show enhanced ability in some respects relative to normal? It is these gross behavioural differences, after all, that led to the initial assumption of 'spared' and 'impaired' (normal) modular capacities in these atypical populations. Hence, while neuroconstructivists may be right to argue that gross behavioural measures are sometimes less indicative than subtle behavioural measures of underlying cognitive structures, they need to give some account of how some underlying similarities in cognitive architecture could nevertheless produce gross behavioural differences in the target populations.

While answering the particular questions about WS and ASD must be complicated in detail, requiring a great deal of empirical research, the neuroconstructivist view of cognitive development provides a theoretical framework within which such questions can be profitably addressed. As mentioned above, neuroconstructivists reject the thesis that cognitive development consists in the predetermined epigenesis of a range of modular capacities. Instead, they regard such development as a probabilistic process that is multiply and interactively constrained by an array of genetically determined neural processing biases in the presence of (developmentally changing) experiential input. Thus, while cognitive development is a matter of increasing cortical specialisation (what many would call 'modularity'), the direction and even degree of such specialisation will depend on at least the following: the sort of (neurally driven) experiential input an individual seeks (e.g. attention to versus avoidance of faces); how efficiently and in what manner that experiential input is processed (a matter of currently existing neuronal structures with their particular response properties or biases); and finally (genetically and experientially determined) architectural constraints on neural growth and development (Karmiloff-Smith 1992, 1998; Johnson 2000; Oliver et al. 2000; Quartz 2003). While some subset of these factors may be shared across developmental disorders—e.g., in WS and ASD, certain biases in the way visual stimuli are processed—many others will not, explaining dramatic differences in phenotypic outcome.

On this issue of face processing in WS and ASD in particular, adopting a neuroconstructivist perspective allows theorists to reconcile data that otherwise seems to be in conflict. On the one hand, WS individuals are unquestionably proficient in many face-processing tasks, certainly relative to ASD individuals and also relative to IQ-matched controls, making them relative 'experts' in this domain. On the other hand, they show the typical neural markers and subtle behavioural properties of a cognitive processing style that shares much in common with that of ASD individuals, who are unquestionably deficient in this domain. Still, because individuals with WS are especially drawn to faces in contrast to other objects (the reverse of what is found in ASD), it seems likely that significantly more of their neural resources become devoted to face-processing, certainly in comparison with ASD individuals, and possibly even in comparison with individuals who develop normally. Consequently, even though the architectural properties of this WS neural system may be quite different from normal—e.g. devoted more to exemplar learning as opposed to the more efficient, higher level extraction of configural information (a feature which turns up in more subtle behavioural measures)—less is sacrificed in the way of gross behavioural performance than might be expected on the basis of IQ or other cognitive measures (Karmiloff-Smith 1997). In short, if this is 'expertise', it is expertise with a different cognitive profile from that of normally developing individuals.

I have discussed the example of face processing in these different populations at some length in order to demonstrate the difference that neuroscientific evidence can make, both to the interpretation of existing behavioural studies and to the design of further studies aimed at generating more refined data to resolve questions about cognitive structure that arise precisely in light of particular neuroscientific findings. The example is also intended to illustrate just how misleading behavioural studies can be if they are taken to stand on their own. At the very least, this suggests that it would be

unwise to embrace the methodological precepts of ultra-CNP, particularly if there is good reason to think that the case of face-processing is representative of a more general phenomenon—that phenomenon being *cognitive* multiple realisability (discussed in Sect. 3).

Neuroconstructivists give us concrete illustrations of this abstract philosophical possibility. Using a variety of examples drawn from other cognitive domains and involving other cognitive disorders, they argue persuasively that naturally developing systems will evolve functionally distinctive solutions to the problem space of any cognitive domain, if the parameters that initiate and shape developmental processes themselves (genetic and environmental) are set in different enough ways (Karmiloff-Smith et al. 2003; Thomas and Karmiloff-Smith 2002; Johnson et al. 2002). Differences in functional specificity will, of course, reveal themselves in different patterns of behaviour; but since the idea of cognitive multiple realisability is that one and the same cognitive task can be solved in various different ways, using different algorithms, the differences in behaviour are likely to be subtle and diffuse rather than crudely obvious. All of this goes to reinforce the point that behavioural evidence alone is likely to be insufficient to distinguish such different functional possibilities. Of course, neural evidence alone—which indicates, for instance, that some cognitive tasks, such as face recognition, are performed using different neural substrates by certain populations—might indicate that the self-same cognitive architecture is realised in physically different ways (*physical* multiple realisability) or it might indicate that the cognitive architecture itself is quite different (*cognitive* multiple realisability). In the end, the only way to disentangle such possibilities is by inference to the best explanation. However, the key point remains that such inferences will be most compelling when they draw on mutually reinforcing evidence from multiple levels of analysis.

Apart from this methodological lesson, there is a deeper, more substantive conclusion to be drawn from the discussion in this section. As argued in previous sections, the methodological stance advocated by proponents of ultra-CNP has no a priori standing but is rather conditional on the acceptability of three substantial ssumptions regarding the structure and development of both normal and atypical cognitive systems. To review, these assumptions include: (i) universality: all human brains are qualitatively similar in terms of an underlying modular cognitive structure; (ii) fractionation or subtractivity: such modular structure can be differentially impaired, either as a consequence of injury to normally developed brains or through abnormal development; and (iii) transparency: when such damage or malformation occurs, it can occur 'locally' in cognitive terms, meaning that other parts of the system will continue to function 'normally' or as they would had there been no damage to, or malformation of, the affected modules. Such assumptions license the view that 'proximal' causes of atypical cognitive functioning (damage to particular hypothesised modules) can be read off behavioural data no matter what the aetiology of the atypical functioning (i.e. no matter what the "distal" cause), because of overall parity in cognitive 'endstates'.

But what is the status of these assumptions? Clearly all sides should agree that they are empirical assumptions, open to some kind of evidential confirmation or refutation. However, proponents of ultra-CNP seem committed to the view that neuroscientific evidence either currently does not—or, more strongly, cannot—have any bearing on

their acceptability. This would explain both their continued insistence on the relative insignificance of neuroscientific evidence to the development of cognitive theory (discussed in Sect. 3) and their seeming acceptance of the circularities involved in neuropsychological inferences as just a feature of regular scientific practice (discussed in Sect. 2). But neuroconstructivists effectively dispute this position. By their arguments, a growing body of work on the kind of processes involved in cortical development and specialisation show the background assumptions of ultra-cognitive-CNP to be inapplicable in the context of developmental disorders.

At the very least, this would suggest a modest conclusion: that the behavioural methods of ultra-CNP are inappropriate to the study of abnormally developed cognitive architectures, since the dissociationist logic on which they rely is vitiated by the failure of their background assumptions (Bishop 1997). However, there is also a stronger conclusion in the offing. If the developmental case provides a demonstration that patterns of 'spared' and 'impaired' behavioural performance can be quite deceptive about the underlying modular architecture of the abnormally developed brain, then this should underscore the fact that patterns of spared and impaired behavioural performance in the case of acquired disorders may be equally deceptive for the purpose of making inferences about the underlying modular structure of normal cognition. Is there any reason to think this might be so?

As we have seen, the idealising assumptions that guide such inferences—universality, fractionation and transparency—are not invulnerable to neuroscientific evidence; and, on this basis, theorists have begun to question their applicability even in cases of acquired pathology (Kosslyn and Van Kleeck 1990; Johnson et al. 2002). More deeply still, the foundational assumption of modularity has itself come under pressure, both from work in connectionist modelling (Plaut 1995; Oliver et al. 2000) and from continuing work in the neurosciences. For instance, with respect to the latter, cross-cutting results from a plethora of lesion and imaging studies raise continuing concerns that cognitive theorists cannot be getting the story right about the modular processes supposedly involved in various cognitive tasks. And though it is now standard practice to explain away these difficulties by pointing out that modular processes need not be anatomically localised, such a rejoinder does not do much to help bring cognitive theory into productive contact with neuroscientific evidence. More radical thinking is needed. One possibility is to take a more bottom-up approach to understanding structure-function correspondences. For instance, Price and Friston (2005) argue that, since many neuronal structures can perform multiple functions depending on other areas with which they interact, a more useful functional taxonomy for such structures would involve characterisations that respect their actual activation patterns across a wide variety of cognitive tasks, even if such characterisations do not correspond to any component processes as these are more traditionally identified through task-specific behavioural studies. Such work will be highly labour intensive, requiring the pooling of results from multiple imaging studies in a variety of cognitive domains. It may also be unpopular, since it will require a radical revision in how cognitive theorists think about 'modular decomposition'. Still, as Price and Friston point out, the ultimate payoff is to develop a taxonomy of functional characterisations that "will enable cognitive and anatomical models to converge and be built in a mutually consistent way" (p. 269).

5 Conclusion

In these and other ways, the more traditional methods of cognitive psychology are coming under pressure, as a burgeoning understanding of the brain encourages theorists to take on board what Daniel Dennett has always warned about: that "the decomposition of one's competence model into parts, phases, states, steps or whatever *need* shed no light at all on the decomposition of any mechanical parts, phases, states of steps of the system being modelled" (Dennett 1987, pp. 76–77). How will the specialised subdiscipline of CNP transform itself under this revolutionising trend? That remains to be seen. But one thing is certain: the methodology and assumptions of ultra-CNP look to be quite unsustainable, despite the best effort of its proponents to argue to the contrary. As Trevor Harley remarks, "the future lies in a synthesis—a combination of computational modelling, analysis of single-case studies, case-series approaches, real-time data, and information from the neurosciences" (Harley 2004b, p. 55). In a word, there is nothing to lose and everything to gain by theorists insisting on the 'neuro' in cognitive *neuro*psychology.

References

Adolphs, R., Sears L., & Piven, J. (2001). Abnormal processing of social information from faces in autism. *Journal of Cognitive Neuroscience, 13*(2), 232.

Baron-Cohen, S. (1989). The autistic child's theory of mind: A case of specific developmental delay. *Journal of Child Psychology and Psychiatry, 30*, 285–298.

Baron-Cohen, S. (1998). Modularity in developmental cognitive neuropsychology: Evidence from autism and Gilles de la Tourette syndrome. In J. Burack (Ed.), *Handbook of mental retardation and development* (pp. 334–348). Cambridge: Cambridge University Press.

Baron-Cohen, S. W. S. J. T. (1997). Is there a "language of the eyes"? Evidence from normal adults, and adults with autism or Asperger syndrome. *Visual Cognition, 4*, 311.

Bellugi, U., Birhle, A., Jernigan, T., Trauner, D., & Doherty, S. (1990). Neuropsychological, neurological, and neuroanatomical profile of Williams syndrome. *American Journal of Medical Genetics, 6*, 115–125.

Bellugi, U., Lechtenberger, W., Jones, W., Lai, Z., & St. George, M. (2000). The neurocognitive profile of Williams syndrome: A complex pattern of strengths and weaknesses. In U. Bellugi & M. St. George (Eds.), *Journey from cognition to brain to gene: Perspectives from Williams syndrome*. Cambridge, MA: MIT Press.

Bellugi, U., Lichtenberger, L., Mills, D., Galaburda, A., & Korenberg, J. R. (1999). Bridging cognition, the brain and molecular genetics: Evidence from Williams syndrome. *Trends in Neurosciences, 22*, 197–207.

Bellugi, U., Wang P. P., & Jernigan, T. (1994). Williams syndrome: An unusual neuropsychological profile. In S. Broman & J. Grafman (Eds.), *Atypcial cognitive deficits in developmental disorders: Implications for brain function*. Erlbaum.

Bishop, D. V. M. (1997). Cognitive neuropsychology and developmental disorders: uncomfortable bedfellows. *Quarterly Journal of Experimental Psychology, 50A*(4), 899–923.

Block, N. (1995). The mind as the software of the brain. In E. Smith & D. Osherson (Eds.), *Thinking: an invitation to cognitive science*. Cambridge, MA: MIT Press.

Boucher, J., & Lewis, V. (1992). Unfamiliar face recognition in relatively able autistic children. *Journal of Child Psychology and Psychiatry, 33*(5), 843–859.

Boucher, J., Lewis, V., Collis, G. (2000). Familiar face and voice matching and recognition in children with autism. *The Journal of Child Psychology and Psychiatry and Allied Disciplines, 39*(02), 171–181.

Caramazza, A. (1986). On drawing inferences about the structure of normal cognitive systems from an analysis of patterns of impaired performance: The case for single-patient studies. *Brain and Cognition, 5*, 41–66.

Caramazza, A. (1992). Is cognitive neuropsychology possible? *Journal of Cognitive Neuroscience, 4*(1), 80–95.
Caramazza, A., & Coltheart, M. (2006). *Cognitive neuropsychology* twenty years on. *Cognitive Neuropsychology, 23*(1), 3–12.
Carey, S., & Diamond, R. (1977). From piecemeal to configural representation of faces. *Science, 195*, 312–314.
Chung, M.-S., & Thompson, D. M. (1995). Development of face recognition. *British Journal of Psychology, 86*(1), 55.
Churchland, P. M., & Churchland, P. S. (1996). Replies from the Churchlands. In R. N. McCauley (Ed.), *The Churchlands and their critics*. Oxford: Blackwell Publishers.
Coltheart, M. (2002a). Cognitive neuropsychology. In J. Wixted (Ed.), *Stevens' handbook of experimental psychology* (vol. 4, pp. 139–174). New York: Wiley & Sons.
Coltheart, M. (2002b). *The distinction between proximal and distal causes of developmental disorders of cognition.* Presented in the Seminar Series, Research School of Social Sciences, Australian National University, Canberra, Australia.
Coltheart, M. (2004). Brain imaging, connectionism, and cognitive neuropsychology. *Cognitive Neuropsychology, 21*(1), 21–25.
Coltheart, M., Curtis, B., Atkins, P., & Haller, M. (1993). Models of reading aloud: Dual route and parallel-distributed processing approaches. *Psychological Review, 100*, 589–608.
Coltheart, M., & Jackson, N. E. (1998). Defining dyslexia. *Child Psychology and Psychiatry Review, 3*(01), 12–16.
Coltheart, M., & Langdon, R. (1998). Autism, modularity and levels of explanation in cognitive science. *Mind and Language, 13*(1), 138–152.
Dalton, K., Nacewicz, B., Johnstone, T., Schaefer, H., Gernsbacher, M. A., Goldsmith, H. H., Alexander, A., & Davidson, R. (2005). Gaze fixation and the neural circuitry of face processing in autism. *Nature Neuroscience, 8*(4), 519–526.
Davies, S., Bishop, D., Manstead, A. S. R., & Tantam, D. (1994). Face perception in children with autism and Asperger's syndrome. *Journal of Child Psychology and Psychiatry, 35*(6), 1033–1057.
Dawson, G., Carver, L., Meltzoff, A. N., Panagiotides, H., McPartland, J., & Webb, S. J. (2002). Neural correlates of face and object recognition in young children with autism spectrum disorder, developmental delay, and typical development. *Child Development, 73*(3), 700–717.
de Hahn, M. (2001). The neuropsychology of face processing during infancy and childhood. In C. A. Nelson & M. Luciana (Eds.), *Handbook of developmental cognitive neuroscience* (pp. 381–398). Cambridge, MA: MIT Press.
de Schonen, S., Mathivet, E., & Deurelle, C. (1989). A timing puzzle. *Current Psychology of Cognition, 9*, 147–161.
Dennett, D. (1987). *The intentional stance*. Cambridge, MA: MIT Press.
Deruelle, C., Mancini, J., Livet, M. O., Casse-Perrot, C., & de Schonen, S. (1999). Configural and local processing of faces in children with Williams Syndrome. *Brain and Cognition, 41*(3), 276–298.
Elgar, K., & Campbell, R. (2001). Annotation: The cognitive neuroscience of face recognition: Implications for developmental disorders. *Journal of Child Psychology and Psychiatry, 42*(6), 705–717.
Ellis, H. D., Ellis, D. M., Fraser, W., & Deb, S. (1994). A preliminary study of right hemisphere cognitive deficits and impaired social judgments among young people with Asperger syndrome. *European Child & Adolescent Psychiatry, 3*(4), 255–266.
Fodor, J. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.
Frith, U., & Happé, F. (1994). Autism: beyond 'theory of mind'. *Cognition, 50*, 115–132.
Gauthier, I., & Logothetis, N. K. (2000). Is face recognition not so special after all? *Cognitive Neuropsychology, 17*(1), 125–142.
Gerrans, P. (2003a). Nativism and neuroconstructivism in the explanation of Williams syndrome. *Biology and Philosophy, 18*(1), 41–52.
Gerrans, P. (2003b). Nativism, neuroconstructivism, and developmental disorder. *Behavioral and Brain Sciences, 25*(06), 757–758.
Gopnik, M. (1997). Language deficits and genetic factors. *Trends in Cognitive Science, 1*, 5–9.
Grice, S., Spratling, M. W., Karmiloff-Smith, A., Halit, H., Csibra, G., De Haan, M., & Johnson, M. H. (2001). Disordered visual processing and oscillatory brain activity in autism and Williams syndrome. *Neuroreport, 12*(12), 2697–2700.

Grill-Spector, K., Knouf, N., & Kanwisher, N. (2004). The fusiform face area subserves face perception, not generic within-category identification. *Nature Neuroscience, 7*(5), 555–562.

Hadjikhani, N., Joseph, R., Snyder, J., Chabris, C., Clark, J., Steele, S., McGarth, L., Vangel, M., Aharon, I., Feczko, E., Harris, G., & Tager-Flusberg, H. (2004). Activation of the fusiform gyrus when individuals with autism spectrum disorder view faces. *NeuroImage, 22*, 1141–1150.

Happé, F. (1999). Autism: Cognitive style or cognitive deficit. *Trends in Cognitive Science, 3*(6), 216–222.

Harley, T. A. (2004a). Does cognitive neuropsychology have a future? *Cognitive Neuropsychology, 21*(1), 3–16.

Harley, T. A. (2004b). Promises, promises. *Cognitive Neuropsychology, 21*(1), 51–56.

Haxby, J. V., Hoffman, E. A., & Gobbini, M. I. (2000). The distributed human neural system for face perception. *Trends in Cognitive Sciences, 4*, 223–233.

Hobson, R. P., Ouston, J., & Lee, A. (1988). Emotion recognition in autism: Coordinating faces and voices. *Psychological Medicine, 18*, 911–923.

Jackson, N. E., & Coltheart, M. (2001). *Routes to reading success and failure: Toward an integrated cognitive psychology of atypical reading*. Psychology Press.

Johnson, M. H. (2000). Functional brain development in infants: Elements of an interactive specialization framework. *Child Development, 71*(1), 75–81.

Johnson, M. H., Halit, H., Grice, S. J., & Karmiloff-Smith, A. (2002). Neuroimaging of typical and atypical development: A perspective from multiple levels of analysis. *Development and Psychopathology, 14*, 521–536.

Joseph, R. M., & Tanaka, J. (2003). Holistic and part-based recognition in children with autism. *Journal of Child Psychology & Psychiatry & Allied Disciplines, 44*(4), 529.

Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience, 17*(11), 4302–4311.

Karmiloff-Smith, A. (1992). *Beyond modularity: A developmental perspective on cognitive science*. Cambridge, MA: Bradford/MIT Press.

Karmiloff-Smith, A. (1997). Crucial differences between developmental cognitive neuroscience and adult neuropsychology. *Developmental Neuropsychology, 13*, 513–524.

Karmiloff-Smith, A. (1998). Development itself is the key to understanding developmental disorders. *Trends in Cognitive Science, 2*(10), 389–398.

Karmiloff-Smith, A., Brown, J. H., Grice, S., & Paterson, S. (2003). Dethroning the myth: Cognitive dissociations and innate modularity in Williams syndrome. *Developmental Neuropsychology, 23*(1/2), 227–242.

Karmiloff-Smith, A., Thomas, M., Annaz, D., Humphreys, K., Ewing, S., Brace, N., Duuren, M., Pike, G., Grice, S., & Campbell, R. (2004). Exploring the Williams syndrome face-processing debate: The importance of building developmental trajectories. *Journal of Child Psychology and Psychiatry, 45*(7), 1258–1274.

Klin, A. (2000). Attributing social meaning to ambiguous visual stimuli in higher-functioning autism and Asperger syndrome: The social attribution task. *The Journal of Child Psychology and Psychiatry and Allied Disciplines, 41*(07), 831–846.

Klin, A., Sparrow, S. S., de Bildt, A., Cicchetti, D. V., Cohen, D. J., & Volkmar, F. R. (1999). A normed study of face recognition in autism and related disorders. *Journal of Autism and Developmental Disorders, 29*(6), 499–508.

Kosslyn, S. M., & Van Kleeck, M. (1990). Broken brains and normal minds: Why humpty dumpty needs a skeleton. In E. Schwartz (Ed.), *Computational neuroscience* (pp. 390–402). Cambridge, MA: MIT Press.

Laing, E., Butterworth, G., Ansari, D., Gsodl, M., Longhi, E., Panagiotaki, G., Paterson, S., & Karmiloff-Smith, A. (2002). Atypical development of language and social communication in toddlers with Williams syndrome. *Developmental Science, 5*(2), 233–246.

Lakatos, I. (1974). Falsification and the methodology of scientific research programs. In I. Lakatos & A. Musgrave (Eds.), *Criticism and the growth of knowledge* (pp. 91–196). Cambridge: Cambridge University Press.

Leslie, A. (1991). The theory of mind impairment in autism: Evidence for a modular Mechanism of development? In A. Whiten (Ed.), *Natural theories of mind*. Oxford: Blackwell.

Mehler, J., Morton, J., & Jusczyk, P. W. (1984). On reducing language to biology. *Cognitive Neuropsychology, 1*, 83–116.

Mervis, C. B., & Bertrand, J. (1997). Developmental relations between cognition and language: Evidence from Williams syndrome. In L. B. Adamson & M. A. Romski (Eds.), *Research on communication*

and language disorders: Contributions to theories of language development (pp. 75–106). New York: Brookes.

Mills, D. L., Alvarez, T. D., St. George, M., Appelbaum, L. G., Bellugi, U., & Neville, H. (2000). Electrophysiological studies of face processing in Williams syndrome. *Journal of Cognitive Neuroscience, 12*(1), 47–64.

Morton, J., & Johnson, M. H. (1991). CONSPEC and CONLEARN: A two-process theory of infant face recognition. *Psychological Review, 98*(2), 164–181.

Oliver, A., Johnson, M. H., Karmiloff-Smith, A., & Pennington, B. (2000). Deviations in the emergence of representations: A neuroconstructivist framework for analysing developmental disorders. *Developmental Science, 3*(1), 1–23.

Osterling, J., & Dawson, G. (1994). Early recognition of children with autism: A study of first birthday home videos. *Journal of Autism and Developmental Disorders, 24*, 247–257.

Passarotti, A. M., Paul, B. M., Bussiere, J. R., Buxton, R. B., Wong, E. C., & Stiles, J. (2003). The development of face and location processing: An fMRI study. *Developmental Science, 6*(1), 100–117.

Pierce, K., Muller, R. A., Ambrose, J., Allen, G., & Courchesne, E. (2001). Face processing occurs outside the fusiform 'face area' in autism: Evidence from functional MRI. *Brain, 124*(10), 2059–2073.

Pinker, S. (1999). *Words and rules*. New York/London: Basic Books/Weidenfield & Nicholson.

Plaut, D. C. (1995). Double dissociation without modularity: Evidence from connectionist neuropsychology. *Journal of Clinical and Experimental Neuropsychology, 17*(2), 291–321.

Price, C. J., & Friston, K. J. (2005). Functional ontologies for cognition: The systematic definition of structure and function. *Cognitive Neuropsychology, 22*(3), 262–275.

Quartz, S. R. (2003). Innateness and the brain. *Biology and Philosophy, 18*(1), 13–40.

Rae, C., Karmiloff-Smith, A., Lee, M. A., Dixon, R. M., Grant J., Blamire, A. M., Thompson, C. H., Styles, P., & Radda, G. K. (1998). Brain biochemistry in Williams syndrome: Evidence for a role of the cerebellum in cognition? *Neurology, 51*(1), 33–40.

Schultz, R. T., Gautier, I., Klin, A., Fulbright, R. K., Anderson, A. W., Volkmar, F., Skudlarski, P., Lacadie, C., Cohen, D. J., & Gore, J. C. (2000). Abnormal ventral temporal cortical activity during face discrimination among individuals with autism and Asperger syndrome. *Archives of General Psychiatry, 57*, 331–340.

Smith, N. (1999). *Chomsky: Ideas and ideals*. Cambridge, UK: Cambridge University Press.

Stone, T., & Davies, M. (1999). Autonomous psychology and the moderate neuron doctrine. *Behavioral and Brain Sciences, 22*(05), 849–850.

Swettenham, J., Baron-Cohen, S., Charman, T., Cox, A., Baird, G., Drew, A., Rees, L., & Wheelwright, S. (1998). The frequency and distribution of spontaneous attention shifts between social and nonsocial, stimuli in autistic, typically developing, and nonautistic developmentally delayed infants. *Journal of Child Psychology and Psychiatry, 39*, 747–753.

Tarr, M. J., & Gauthier, I. (2000). FFA: A flexible fusiform area for subordinate-level visual processing automatized by expertise. *Nature Neuroscience, 3*(8), 764–769.

Temple, C. M. (1997). *Developmental cognitive neuropsychology*. Hove, UK: Psychology Press.

Thomas, M., & Karmiloff-Smith, A. (2002). Are developmental disorders like cases of adult brain damage? Implications from connectionist modelling. *Behavioral and Brain Sciences, 25*(06), 727–750.

Thomas, M., & Karmiloff-Smith, A. (2003). Residual normality: Friend or foe? *Behavioral and Brain Sciences, 25*(06), 772–780.

Udwin, O., & Yule, W. (1991). A cognitive and behavioral phenotype in Williams syndrome. *Journal of Clinical and Experimental Neuropsychology, 13*, 232–244.

Van Orden, G. C., Pennington, B. F., & Stone, G. O. (2001). What do double dissociations prove? *Cognitive Science, 25*(1), 111–172.

Wang, P. P., Doherty, S., Rourke, S. B., & Bellugi, U. (1995). Unique profile of visuoperceptual skills in a genetic syndrome. *Brain and Cognition, 29*, 54–65.