Thinking through the implications of neural reuse for the additive factors method

Abstract

One method for uncovering the subprocesses of mental processes is the "Additive Factors Method" (AFM). The AFM uses reaction time data from factorial experiments to infer the presence of separate processing stages. This paper investigates the conceptual status of the AFM. It argues that one of the AFM's underlying assumptions is problematic in light of recent developments in cognitive neuroscience. Discussion begins by laying out the basic logic of the AFM, followed by an analysis of the challenge presented by neural reuse. Following this, implications are analysed and avenues of response considered. Keywords: additive factors method; seriality assumption; anatomical modularity; neural reuse.

Keywords: additive factors method, neural reuse, stage models, seriality assumption

Introduction

A good place to start when trying to understand a complex process or system is to determine its constitutive parts or modules. For example, to figure out how people succeed in visual search during reading, the time between stimulus and response can be broken down into an encoding, feature extraction and identification stage (Resink, 2005; Tovey & Herdman, 2014). The decomposition of the time between the stimulus and response enables discovery of the underlying subprocesses. The stimulus–response time intervals reflect the series of processing stages underlying complex behaviours.

One method for uncovering the subprocesses of mental processes is the "Additive Factors Method" (henceforth AFM) (Townsend & Nozawa 1995; 2001, 2011, 2013; Coltheart, 2011). The AFM uses reaction time data from factorial experiments to infer the presence of separate modules or processing stages. A mental process can be broken down into subprocesses when those subprocesses are 'separately modifiable' – that is, when each of the proposed modules can be modified without effect to the other, and vice versa. For example, to show that two stages A and B are separately modifiable it must be shown that two factors, F and G, affect only either A or B, but not both. In other words, F can affect A and G can affect B, but not the reverse. The result of an AFM analysis is what are called 'stage models'.

This paper investigates the conceptual status of the AFM. It argues that one of the AFM's main assumptions is problematic in light of recent developments in cognitive neuroscience. In particular, the argument is that theories of neural reuse present a challenge to the conceptual link between AFM's 'seriality assumption' and the single processor cases it relies on. Discussion begins by laying out the basic logic of AFM, followed by an analysis of the challenge presented by neural reuse theories. Implications are then analysed and avenues of response considered.

The Additive Factors Method

Factorial experiments are studies in which the effects of two or more variables are investigated by manipulating the presence of each factor across various conditions. In its simplest version (the complete factorial experiment), two factors are studied by comparing the difference each factor has on some measure of performance, such as reaction time. For example, to evaluate the effect of familiarity on pattern recognition, orientation (the rotation of a pattern) can be compared to familiarity (whether subjects are better or worse at recognising the pattern) (Tovey & Herdman, 2014). If orientation has an effect on familiarity, then conditions in which stimuli are presented with different orientations, e.g. 0^0 vs. 90^0, will result in delays in the time required to recognize a pattern.

Factorial experiments form the raw data of the AFM. Factorial data indicates whether two or more factors have either an additive or interaction effect on mean reaction time. Leaving interaction effects to one side for the moment, an additive effect involves two or more factors selectively influencing individual stages of a process. So, for example, if stage A normally takes 10ms and stage B normally takes 15ms and F influences the length of A by 5ms and G influences the length of B by 7ms, then the total duration of time to complete the process that includes stages A and B will be the result of the presence of F and G. The total duration of a process is simply the added the sum of each stage as influenced by each factor.

Factorial experiments supply modifiability information by revealing the selective influence of some factor(s) (Miller et al., 1995). When two or more factors affect the total duration of a process (measured using reaction time), the process can be separated into different modules or processing stages. When patterns of factor effects are observed, a set of hypothesized stages and factor relations that underlie the pattern are proposed. The effects inferred from the factorial experiments are what support inferences about the processing stages, justifying the decomposition of a process or stimulus-response interval into distinct subprocesses.

The Seriality Assumption

One key assumption of the AFM is what Sternberg (2001, 2013) calls the 'seriality' assumption. The seriality assumption says that the AFM can only be applied to processes that are sequentially arranged. For a process to be sequentially arranged, one of two situations must hold, either:

(i) the process must be data-dependent or (ii) a single processor must be responsible for carrying out the process.

In the first set of cases, seriality depends on information being passed along from a previous stage of the process. To use a simple example, heading home from grocery shopping (stage 2) requires first having collected and bagged the groceries at the store (stage 1). In the second set of cases, seriality depends on a process being the result of a single processor. So, for example, if one bakes with two hands multiple steps can be accomplished in parallel, e.g., cracking and whisking eggs; while if one bakes with only one hand, then the process is limited to being complete one task at a time, e.g., cracking each egg individually and then whisking them all together.

How a single process relates to a given processor or set of processors can also vary considerably from case to case. For example, for even a three stage process, there are several types of relations that might hold: (i) a separate processor might carry out each process, (ii) the same processor might carry out every process, or (iii) there be might some combination of the two, where one processor carries out two processes and another processor carries out one process. Figure 1 provides an illustration.

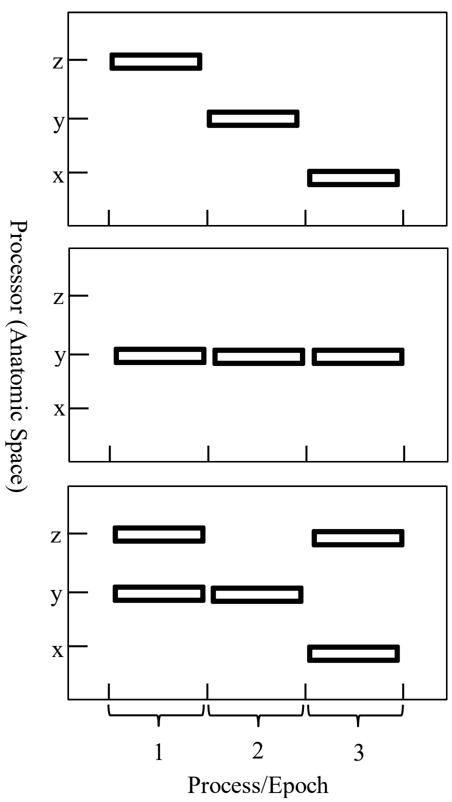

Figure 1. Possible relations between processors and processes for a three stage process.

What is interesting is that the seriality assumption maintains that in at least a subset of cases specialised structures are responsible for carrying out mental processes. That it is possible to find one-to-one mappings between process modules or stages and processing devices. Sternberg (2001), for instance, notes: "Perhaps more surprising is the finding of operations that are partially or wholly sequential when there is no data-dependence…The basis for the sequential structure in such cases may be that the system that carries out the set of operations, possibly the same single processor, is inherently limited in capacity" (original italics, p.735). Not only can processes be sequentially arrange when they involve data dependence but they can also be sequentially arranged when the realising processor has a limited capacity.

However, notice that this is a claim about neural organisation. The single processor view says that neural organisation will take a 'modular' form in certain cases. That in at least some cases mental processes are implemented or realised by dedicated pieces of neural hardware. The claim is that what makes it possible for a given process to be serially arranged is the physical constraints of the realising processor. The view is one of a neural organisation wherein a particular sequential process is carried out by a chainlike structure of connected processing units. To support this claim, for example, Sternberg (2001) appeals to cases of highly specialised anatomical structures, such as the visual cortex of Macaque monkeys. Call this 'anatomical modularity'.

Of course, anatomical modularity is usually considered a 'functional' theory, whereas processing stages are periods of time. Sternberg (2001), for instance, notes: "A stage theory says nothing about the pieces of physical anatomical machinery that carry out the operations in the two stages…information 'transmitted from one stage to the next' does not necessarily go from one place to another; the phrase is unfortunate because it suggests otherwise (p.732). Processor devices that carry out process stages might have functional properties, but the processing stages themselves, at least as informed by the AFM, are neutral with respect to such questions (Kersten, 2016). Nonetheless, there are reasons to see the two views as sides of the same coin. This is because while the processing stages themselves may not have functional properties, they are realised by processors that do, i.e. neurological structures. The point here is simply that the seriality assumption makes specific a claim about the relation between such subprocesses and processors. It does not make a claim about what features those subprocesses have.

It will be worth dwelling on this point as it crucial for the argument to follow. For one might wonder whether 'anatomical modularity' is not better understood as a claim about 'functional' organisation or architecture. If so, then the AFM would be involved in a form of functional decomposition, as it would be set to uncover functional architecture rather than neural organisation.

Crucially, this is not the case. Stage models are set to uncover the subprocesses of mental processes, such as those involved in visual search, understood as epochs or periods of time. Despite some shared inferential machinery, such as factorial experiments, the target and output of the AFM is importantly different from those methods aimed at providing functional decomposition (see Carruthers, 2006).

A brief case study will help flesh out the point further. Consider Tovey and Herdman's (2014) investigation of visual search during reading. Using the effects of familiarity on change perception via a 2 × 5 × 2 factorial design, Tovey and Herdman examined the effects of orientation (upright vs. inverted, set size (4, 7, 10, 13, 16) and change size (Small vs. Large) across four different experiments. In line with Rensink (2005), they suggested that change perception was divided into three process modules: a pre-processing stage, a feature extraction stage, and an identification stage. They proposed that an interaction between change size and orientation and change size and stimulus quality indicated that change size exerted an effect not only on the feature extraction stage but also on the identification stage of change perception. Figure 2 provides an illustration of the model.

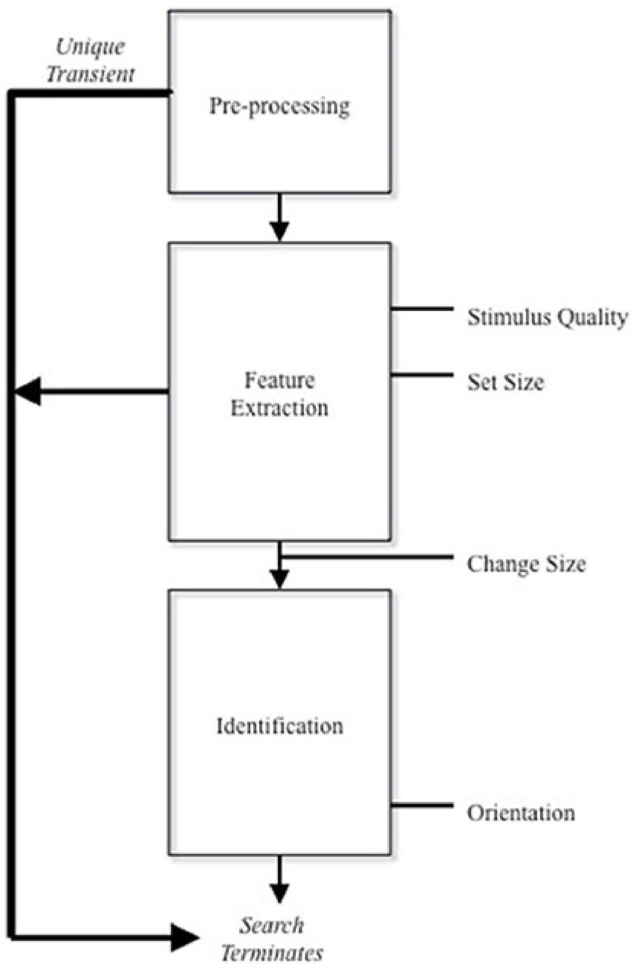

Figure 1. Tovey and Herdman's stage model. Visual search is divided into three stages: a pre-processing stage, a feature extraction stage, and an identification stage.

To explain the effects of change size, Tovey and Herdman proposed a 'gating' mechanism. The gating mechanism redirects information to different stages of the process via detecting changes in size, either by passing the information on to the feature-extraction stage for further processing (assuming the changes are large) or by retaining and verifying the information at the identification stage (assuming the changes are small). The problem is that introduction of a gating mechanism complicates interpretation of Tovey and Herdman's model as a stage model. This is because it introduces functional properties into the model.

Notice that Tovey and Herdman place change size outside of the processing stages, after feature extraction but before identification. This changes the structure of the diagram from a flowchart to a circuit diagram. The arrows no longer represent a succession in time of a series of processes but instead denote the flow of information. The gating mechanism is conceived of as the change size, representing the redirection of information from one stage to another, not only how change size influences time duration.

However, if processing stages are events in time, they need to be strung together end to end, as in a flowchart. If the model represented the effect of change size, it would have to effect the period of time as represented by the box, not the passage or succession of time as represented by the arrows. When represented as a circuit diagram – that is, as describing how processing devices are connected – stage models misleading suggest that the process stages are also processing devices; an interpretation, which, as mentioned, fails to acknowledge the variety of possible relationships that might obtain between process stages and processing devices. Sternberg (2001) frames the point nicely: "It is remarkably easy to slip into a mode of thinking in which stages are processors rather than processes, actors rather than actions; confusion about what a stage might be finds its way into much writing on the subject, even by experts" (p.828). So while Tovey and Herdman's results may be correct, their inclusion of a functional property complicates interpretation of the model (Kersten, 2016).

The ambiguity introduced by the gating mechanism is suggestive of the nature of stage models. For if the AFM is to uncover the subprocesses (understood as epochs of time) of mental processes, then it cannot do so by revealing the functional properties of cognitive systems. If it did, this would blur the distinction between the AFM and other experimental methods.

To illustrate, consider the method of double-dissociation. If two factors F and G are damage to different parts of the brain, and one can show using some measure, such as an EEG, that factor F influences performance on some task A but not task B, while G influences B but not A, then one can infer that the F and G perform different functions. The separate modifiability of tasks A and B by factors F and G on tasks A and B indicate that F and G are have different functions.

Contrast this with the AFM. Whereas double-dissociation uses a direct measure for separate modifiability (the differential activity of different brain regions), the AFM only indirectly tests for separate modifiability via mean-reaction times. It is interested in how a given process can be separated or 'cut' via finding the selective influence of different factors. What this means is that the focus in stage models is on temporal rather than functional properties. Thus, one thing that cannot be meant by the seriality assumption is that what constrains processing stages is functional modularity.

The general point is that anatomical modularity forms more than just a peripheral assumption within the AFM. Indeed, it is what helps, in part, justify inferring the presence of serial processes. If two processes are not data-dependent and yet perform the same function, then it is safe to assume that they are realised by the same processor. That anatomical modularity should underlie part of the seriality assumption is not an insignificant result. The problem is that anatomical modularity is increasingly being called into question.

Neural Reuse and the AFM

Many of the cognitive functions once thought to have dedicated, isolated neural localisations (e.g., Broca's area) are increasingly shown to engage a diverse range of neural units. In a recent meta-review, for example, Anderson and Pessoa (2011) found that 78 different anatomical regions were active in 95 tasks across 9 cognitive domains. Accounts of 'neural reuse' aim to describe how different brain regions often exploit, recycle, or redeploy neural circuitry for various cognitive ends (Hurley 2005; Dehaene & Cohen, 2007; Dehaene, 2009; Anderson 2007, 2011, 2014).

A large swath of evidence now favours neural reuse as a thesis of neural organisation. To spare a long digression, consider a small sampling of some characteristic studies. Glenberg and Kaschak (2002), for instance, show that when asked to make sense judgments about different sentences participants take longer on sentences that run counter to the required action than those that do not. Richardson et al. (2003) show that certain sets of verbs, such as 'hope' and 'respect', activate meaning-specific spatial schemas. Pulvermuller (2005) demonstrates that listening to action words, such as 'lick' or 'pick', activate regions of the primary motor cortex, areas often associated with the actions themselves. Casasanto and Boroditsky (2008) show that people are often unable to ignore irrelevant spatial information when making judgments about duration, but not the converse. That mental representations of time are intimately tied up with perceptions of space. Finally, Casasanto and Dijkstra (2010) demonstrate that there is a bidirectional influence between motor control and autobiographical memory.

Neural reuse theories raise a number of interesting questions about cognition, such as whether a new 'cognitive ontology' needs to be developed (Anderson, 2014). However, for present purposes, the point to note is that neural reuse also raises questions for the second set of cases appealed to by the seriality assumption: namely, that some processes are sequentially arranged in virtue of being realised by single processors. The issue is that if neural reuse is true, then it is unlikely that there will be any single process that has a unique anatomical structure or processor supporting it. Finding a one-to-one mappings between processor and process will prove particularly troublesome if neural regions support multiple operations.

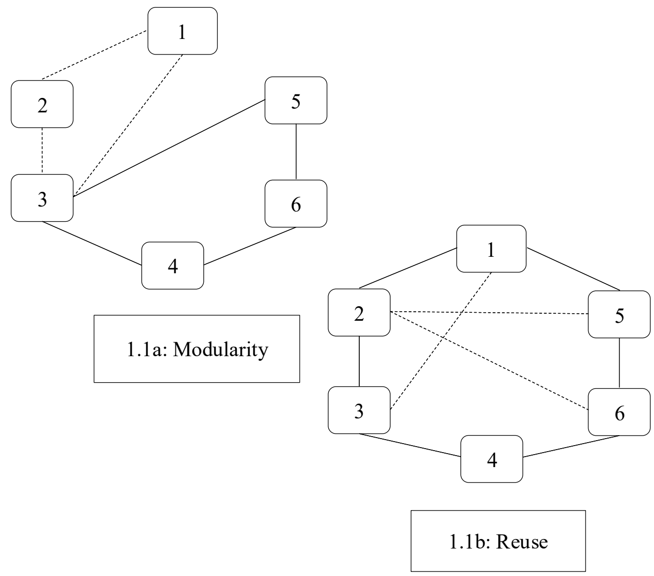

Figure 3.Two possibilities of neural organisation for two cognitive operations.

There seem to be two options. 1.1a represents a modular design, where each cognitive operation has specific dedicated neural circuitry. This is what is required by the second half of the seriality assumption. 1.1b, on the other hand, represents a neural reuse design, where each cognitive process is shared among a number of neural circuits. The AFM requires that 1.1a hold for at least a subset of cases. However, if, as noted, neural reuse is true, then whatever else might be right about the AFM, the single processor cases might not exist. Neural reuse seems to challenge the link between the seriality assumption and one of its inferential bases.

It is important to be clear about this point. For it might be still maintained that the AFM uncovers something about functional organisation. That it would not matter if the same neural hardware were involved in multiple operations because once those operations were fixed the AFM would be set to uncover functional organisation.

But again, once it is appreciated that the target and output of the AFM is not functional models but models of temporal stages it follows that the underlying assumption about processors has to be about neural organisation. For although it is right to point out that neural architecture can stand in complex relations to functional architecture, such as distributed neural regions supporting a functionally modular architecture, such considerations cannot do much work here.

This is because they threaten to undermine the AFM's conceptual standing. If AFM did reveal insights into functional architecture, then its distinctiveness would be undercut, for it would no longer reveal insights into the organisation of cognitive subprocesses understood as temporal sequences.

Another tack would be try to accept the incompatibility of anatomical modularity and neural reuse but nonetheless reject neural reuse on the basis of the wider importance of the AFM. One might argue, in other words, something along the following lines: (1) AFM is essential to psychology; (2) AFM is incompatible with neural reuse; therefore, (3) neural reuse is false.

But there are at least two problems with this type of argument. One is that it assumes that cognitive psychology can operate independently of cognitive neuroscience. That the conceptual autonomy of psychology ensures the survival of the AFM. However, the increasing integration of neurological data into cognitive theorising and modelling makes it unlikely that cognitive psychology will continue to function independently of the findings of cognitive neuroscience in this way (Forstmann et al. 2011; de Hollander et al., 2016). The other is that the argument problematically assumes that the choice facing the proponent of AFM is either/or: that either neural reuse has to be rejected or the AFM does. But no such dichotomy is required. As is argued later, it is possible to endorse a version of the AFM that drops anatomical modularity but which still nonetheless operates in other sets of cases.

Three Options for the AFM

It seems fair to say, then, that neural reuse casts some doubt on the inferential bases of one of the key assumption of the AFM. Given this, three options seem available to the proponent of the AFM. One is to drop the seriality assumption altogether. One might maintain that the AFM can continue on without the seriality assumption. This option seems undesirable insofar as the seriality assumption is part and parcel of the AFM logic. Separate modifiability only makes sense when the processes being investigated are sequentially arranged. Dropping seriality would be tantamount to dropping the method altogether; and scuttling the method altogether seems undesirable given the good deal of fruitful research that has been carried out using the AFM (e.g., Resink, 2005; Tovey & Herdman, 2014).

A second option would be to reform the seriality assumption in light of neural reuse. One might claim, for instance, that serial processes can be the product of distributed neural processors. The problem with this option is that it undermines the inferential link between processor and processing stages. Anatomical modularity forms a key assumption within the AFM. Without it, the AFM would lose its ability to infer a serial ordering. Return, for example, to the baking case, it is only because there is one single processor that the stages are arranged serially. The addition of a second hand opens up the task to being achieved in multiple stages, i.e. in parallel. If multiple processors are admitted, then inferences to processing stages are underdetermined.

But, one might object, it could be that a bunch of miniature 'hands' accomplish the baking task. In other words, that a distributed network of miniature processors performs the task serially, whose actual decomposition is discoverable (at least in part) by the AFM. The problem with this rejoinder is that again misses the key point of stage models, and to lesser extent the baking example. For while it is true that adding more processors speeds up the process, it also makes it impossible to interpret the process as serially ordered. For example, switching to using two hands during the baking (i.e. allowing multiple processors) opens the process up to being completed in parallel. There is nothing that forces the process into being completed in successive stages. Thus, in assuming that a process is realised by distributed set of processors one undermines the ability to interpret that process as serial in the first place. The grounds for inferring seriality rests on the process being carried out by a single processor.

Finally, one might jettison the seriality assumption's commitment to the single processor view, i.e. anatomical modularity. This might preserve what is right about the AFM (i.e. inferring seriality on the basis of data-dependence), while dropping the theoretically suspect part (i.e. reliance on single processor cases). The idea would be to restrict the set of cases under which seriality could be legitimately inferred. That is, whereas previously cases of single processors and data-dependence cases could be used, now only data-dependent cases would be allowed to infer seriality.

For example, a study such as Tovey and Herdman's would not be affected according to the third proposal, because visual search during reading is a data-dependent process. The serially ordering is dependent on each of the previous stages being completed before the next one begins; one cannot, for instance, detect the presence of certain letters before those letters have been registered by the visual system. Tovey and Herdman's study does not rely on the single process cases to work, so it can still be used to infer separate modifiability. However, cases where seriality is inferred because of the supposed presence of a single processor, such as Scarbourgh and Landauer's (1981) study on word repetition effects, would have to be dropped according to this proposal. So, although the removal of anatomical modularity might involve the loss of some of the AFM's methodological punch, as not an insubstantial number of cases involve the assumption (Sternberg, 2001, p.831-2), the method itself would still be preserved in an attenuated form.

Given the spread options, the third proposal seems the most preferable going forward. The first and third options suggest either too high a methodological price or an endorsement of a conceptual tension. Only option three seems to allow the AFM to continue on, though in slightly modified form. On the third proposal, serial stage models can be inferred, but only on the basis of data-dependent cases. The single processor cases underlying the seriality assumption need to be bracketed, at least until such a time that neural reuse can be thoroughly vetted. This might be a welcomed result for some (Stanford & Gurney, 2011; Sternberg, 2013), but maybe not so much for others (Coltheart, 2011).

So, to summarize, though not a devastating blow, neural reuse does represent a serious challenge to some aspects of the AFM. Insofar as neural reuse presents a challenge to anatomical modularity, and anatomical modularity falls out of the seriality assumption, some of the AFM's conceptual foundations need to be reworked. The methodological implications still need to be worked out, but the conceptual moral seems relatively clear: the seriality assumption can no longer rely on single processor cases. Hopefully, then, in having identified the problem and charted some potential responses, the AFM can be put on surer theoretical footing going forward.

References

Anderson, M. (2007). Massive redeployment, exaptation, and the functional integration of cognitive operations. Synthese 159(3): 329–45.

Anderson, M. (2010). Neural reuse: A fundamental organization principles of the brain. Behavioural and Brain Sciences 33(4), 245-266.

Anderson, M. (2014). After Phrenology: Neural Reuse and the Interactive Brain. Cambridge, MA: MIT Press.

Anderson, M., and Pessoa, L. (2011). Quantifying the diversity of neural activations in individual brain regions. In (eds.) L. Carlson, C. Holscher, & T. Shipley, Proceedings of the 33rd Annual Conference of the Cognitive Science Society. Austin, TX: Cognitive Science Society.

Bergeron, V. (2007). Anatomical and functional modularity in cognitive science: Shifting the focus. Philosophical Psychology, 20(2): 175–95.

Carruthers, P. (2006). The Architecture of the Mind: Massive modularity and the flexibility of thought. Clarendon Press/Oxford University Press.

Casasanto, D., and Boroditsky, L. (2008). Time in the mind: Using space to think about time. Cognition, 106, 579-593.

Casasanto, D., and Dijkstar, K. (2010). Motor action and emotional memory. Cognition 115(1), 179-185.

Coltheart, M. (2001). Assumptions and methods in cognitive neuropsychology. In The handbook of cognitive neuropsychology, B. Rapp, ed, 3–21. Psychology Press.

Coltheart, M. (2011). "Methods for modular modelling: additive factors and cognitive neuropsychology." Cognitive Neuropsychology 28, 224–240.

Dehaene, S., and Cohen, L. (2007). Cultural recycling of cortical maps. Neuron 56: 384–98.

Dehaene, S. (2009). Reading in the brain. Viking.

Glenberg, A,, & Kaschak, M. (2002). Grounding language in action. Psychonomic Bulletin & Review, 9, 558-565.

Forstmann, B.U., Wagenmakers, E., Eichele, T. (2011). Reciprocal Relations between Cognitive Neuroscience and Cognitive Models: Opposites Attract? Trends in Cognitive Science, 15(6), 272-279.

Hurley, S. (2005). The shared circuits hypothesis: A unified functional architecture for control, imitation and simulation. In (ed.) Susan Hurley and Nick Chater, Perspectives on imitation: From neuroscience to social science (pp. 76–95). MIT Press.

de Hollander, G., Forstmann, B.U., Brown, S.D. (2016). Different Ways of Linking Behavioural and Neural Data via Computational Cognitive Models. *Biological Psychiatry Cognitive Neuroscience Neuroimaging*, 1(2), 101-109.

Miller, J., van der Ham, F., and Sanders, A. (1995). Overlapping stage models and reaction time additivity: effects of the activation equation. Acta Psychologica (Amst.), 90, 11–28.

Kersten, L. (2016). Processor vs. Processing Accounts of Stage Models: A Cautionary Tale. Frontiers in Psychology, 719, 1-3.

Scarborough, D., and Landauer, T. (1981). Lexical decisions about pairs of adjacent words: A reaction time analysis. Bell Laboratory Technical Memorandum TM91-112215. September.

Stafford, T., & Gurney, K. (2011). Additive factors do not imply discrete processing stages: a worked example using models of the Stroop task. Frontiers in Psychology, 2:287, 1-9.

Sternberg, S. (2001). Discovering mental processing stages: The method of additive factors. In (ed.) S. Sternberg, An invitation to cognitive science (pp. 703-863). Cambridge, MA: MIT Press.

Sternberg, S. (2011). Modular processes in the mind and brain. Cognitive Neuropsychology, 28, 156-208.

Sternberg, S. (2013). The meaning of additive reaction-time effects: some misconceptions. Frontiers in Psychology 744(4), 1-3.

Pulvermuller, F. (2005). Brain mechanisms linking language and action. Nature Reviews Neuroscience, 6, 576-582.

Tovey, M., and Herdmen, C. (2014). Seeing changes: How familiarity alters out perception of change. Visual Cognition, 22(2), 214–238.

Townsend, J., and Nozawa, G. (1995). Spatio-temporal properties of elementary perception: An investigation of parallel, serial and coactive theories. Journal of Mathematical Psychology, 39, 321–360.

Rensink, R. (2000). Seeing, sensing, and scrutinizing. Visual Resolution, 40, 1469–1487.

Rensink, Ronald. (2002). Change detection. Annual Review of Psychology, 53, 245–277.

Rensink, R. (2005). Change blindness. In (ed.) L. Rees and J. Tsotsos, Neurobiology of Attention (pp.76–81). San Diego, CA: Elsevier.