Available online at www.sciencedirect.com
ScienceDirect
Current Opinion in Behavioral Sciences

On the necessity of integrating multiple levels of abstraction in a single computational framework

Hedderik van Rijn[1], Jelmer Borst[2], Niels Taatgen[2] and Leendert van Maanen[3]

We argue that it is imperative that modelers select the right, and potentially differing levels of abstraction for different components of their computational models, as too global or too specific components will hinder scientific progress. We describe ACT-R, from the perspective that is a useful modeling architecture to support this process, and provide two examples in which mixing different levels of abstraction has provided us with new insights.

Addresses
[1] Department of Psychology, University of Groningen, The Netherlands
[2] Department of Artificial Intelligence, University of Groningen, The Netherlands
[3] Department of Psychology, University of Amsterdam, The Netherlands

Corresponding author: van Rijn, Hedderik (Hedderik@van-rijn.org)

Current Opinion in Behavioral Sciences 2016, **11**:116–120

This review comes from a themed issue on **Computational modeling**

Edited by **Peter Dayan** and **Daniel Durstewitz**

http://dx.doi.org/10.1016/j.cobeha.2016.07.007

Introduction

What is the goal of a computational model? We build a model to increase our understanding of human cognitive behavior. Recreating the observed phenomena through simulation demonstrates that all important aspects have been accounted for. However, because it is not feasible to simulate cognition up to the finest level of detail, theoretical importance will dictate on which aspects to focus. This means that for each of the components involved in the explanation of cognitive performance, the right level of abstraction has to be selected.

When a computational model is constructed to explain how linguistic reference processing is constrained by memory load (e.g., [1]) equipping that model with a complete computational account of how photons hitting the retina results in the encoding of letters and eventually words is typically considered unnecessary. Instead, such models assume that a visual system parses the external world and provides the core model with useable input chunks (e.g., 'she', or 'the woman'). Based on this input, the model starts applying a sequence of linguistic processing rules (e.g., 'she' refers to the most recently introduced female agent), and stores and recalls information from memory. By comparing the model's output, collected under different memory capacity settings, with empirical data collected from participants under different memory load conditions, evidence for the role of memory capacity in language processes is collected. Although it would, in principle, be possible to extend this model by incorporating detailed models of the visual system (e.g., Leabra, [2•]), such an extension would make the model overly complicated and detract attention — of both modeler and reader — from the main processes studied.

It is not unlikely that the success of computational modeling paradigms that focus on a single aspect of cognitive processing is partly due to exactly this aspect: they allow the modeler and reader to focus on the most important aspects of the task at hand. For example, models that implement drift-diffusion or linear-ballistic accumulator processes focus on those task aspects that influence decision making, and assume that one parameter indexes how much time is spent on all non-decision related processes. Because of this focus and the model's intrinsic properties, these models allow for very precise explanations of the studied tasks. At the same time, the focus of these models also constrains their applicability. For example, when a sequential series of decisions needs to be made, and the outcome of a first decision influences the motor preparation associated with the next decision, one-shot decision making models need complex extensions to account for the phenomena observed in more complex tasks.

Although both cases discussed above describe models where the authors have abstracted away from parsing visual input, other authors have specifically focused on extending models of basic cognitive processing by constructing better accounts of visual processing. An example of such a model is described by Engelmann *et al.* [3•], who demonstrates that extending a language model by incorporating a precise account of eye movements during reading provides new explanations for existing phenomena which can be tested in subsequent work. This is typical for models of more complex tasks, as accounting for the interaction with the environment (both perception and action) is often crucial for explaining the phenomena of interest (e.g., [4,5]). Therefore, although focusing on a single element can improve the understanding and explainability of the modeling effort, it is often necessary to include interactions with the environment in a model to provide a satisfactory explanation of the data. This raises the question how to determine which level of abstraction to aim for: when is the extension of a model warranted given the additional explanatory power?

Often cited as general law [6], it is thus essential for complex computational models to strive for a balance between complexity and simplicity. However, instead of an overall guiding principle, this version of Ockham's razor can also be applied at a more modular level. The key components of a computational model should be accounted for with as much precision as possible, whereas the auxiliary components can be kept as simple as possible. This would, for example, allow for the development of integrated computational neuroscience models were the most critical aspect of the task is modeled at a biologically plausible level, whereas other components are modeled at a level that is typically associated with traditional information-processing models. However, as soon as different models have slightly different goals, for example because the modelers aim for explaining slightly different phenomena, using Ockham's razor to decide between models becomes problematic are no direct comparisons can be made.

Categorizing computational models in terms of different levels of description is hardly new. Marr and Poggio's levels of analysis [7] categorize information processing systems at three distinct, but complementary levels, ranging from the most global computational level, via an algorithmic level, to the physical level. Similarly, Newell's bands of cognition [8] describe cognitive performance at four levels of description; the Social, Rational, Cognitive and Biological Bands (see also [9,10]). Here we argue that modern computational models should not be analyzed at just one of these levels, but that within a single computational model different levels, analogous to the levels or bands proposed earlier, can, and often should, be combined (see also [11]).

In the remainder of this review we will discuss how the ACT-R cognitive architecture can be used as a modular framework in which modules of different levels of abstraction can be combined into one coherent computational model. We will first briefly introduce ACT-R, followed by two examples of how ACT-R's standard modules can be substituted with more detailed accounts of specific processes. These two examples illustrate, respectively, how more detailed accounts can extend a computational models' scope, and how additional modeling constrains can be derived from more precise modeling of a single module.

ACT-R: a modular cognitive architecture

An example of a framework that can be used to integrate modules of different levels of description is the ACT-R cognitive architecture [12]. ACT-R is first and foremost a psychological theory, for instance explaining how our declarative memory system functions (e.g., [13••]). In addition, ACT-R is implemented as a computational framework, in which simulations of particular tasks can be instantiated (e.g., recovering from interruptions, [14]). Thus, the architecture ACT-R provides the general psychological theory — the architecture of the mind — while models run on this architecture to simulate behavior in particular tasks. Importantly, ACT-R was developed to model complete tasks, from perception to action. An ACT-R model typically interacts with the same interface (albeit often simulated) as human participants, providing response times, accuracy measures, and even fMRI predictions that are directly comparable to those of the participants [15–18].

In ACT-R, the human cognitive system is conceived as a set of independent modules that are coordinated by a central production system (Figure 1; see Anderson, 2007 [12], Chapter 2, for an extensive discussion on the modularity of the mind. Note that although ACT-R assumes strict modularity, recent work has clearly indicated that, for example, different areas of the prefrontal cortex have a differential involvement in decision making). Two modules are used to perceive the environment (aural, visual) and two modules can be used to make responses (manual, vocal). Furthermore, several modules process information internally: the control module (instrumental in cognitive control), declarative memory, and the problem state, a module comparable to the focus of attention in theories on working memory [19,20]. Crucially, each module adheres to a set of precisely defined properties, which are instantiated in the computational framework.

A possible drawback of ACT-R is that all modules function at the same (mostly symbolic) level. Sometimes, a more detailed level of description is needed to account for a specific task, even though it is still useful to simulate the other processes (i.e., perception, action) at a fairly abstract level. The basic idea proposed in this paper is that cognitive architectures could be used as general frameworks, and replace the default modules architecture with more precise implementations, to explain detailed patterns in the data while still simulating the complete task.

Integration of ACT-R and accumulator models: RACE/A

The RACE/A extension to ACT-R [21] provides such an integrative computational model. In default ACT-R, the latency of retrieving an item from memory is determined at the moment a retrieval request is specified, and as such is a ballistic process. However, interference studies have shown that if a competing stimulus is presented *during* the

Figure 1

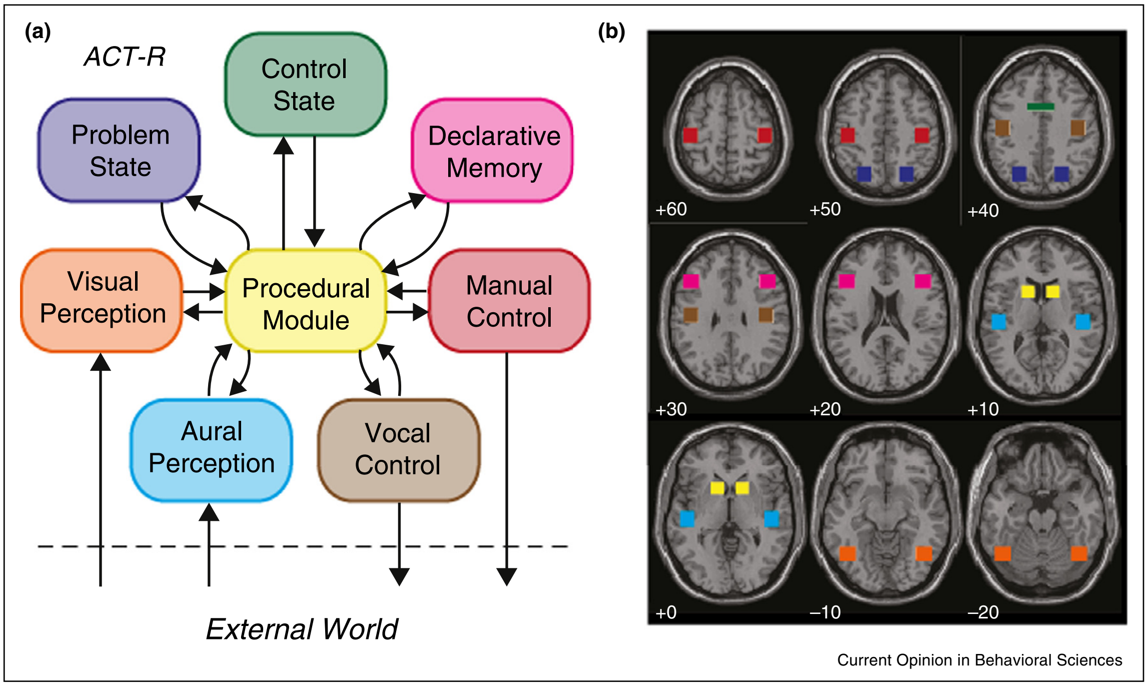

The ACT-R cognitive architecture **(a)** and its mapping to brain regions **(b)**. The colors of the modules correspond to the colored squares in the brain.
Source: Reprinted with permission from Borst and Anderson (2015) [15].

retrieval time, the eventual latency is modulated — falsifying the ballistic nature of the default memory retrieval mechanism. In RACE/A, ACT-R's default declarative memory retrieval processes, modeled as a deterministic choice process (based on the current activation value), are replaced by an accumulator process instantiated as a Leaky Competitive Accumulator model (LCA, [22]). The LCA model assumes that a choice (in this case, a declarative memory retrieval) is the result of sequentially sampling evidence in favor of various choice alternatives (i.e., memory elements). In case of a competition between memory chunks, RACE/A predicts slower retrieval times than standard ACT-R, in accordance with empirical data (e.g., [23]). When there is no competition, RACE/A predicts the same latency as the ACT-R declarative retrieval model [21].

This integration of a lower-level model in a higher level framework demonstrates the feasibility of this approach. When modelers are interested in the low-level intricacies of memory interference processes in cognitive performance, they can integrate the RACE/A extension into their models, but is memory interference is not relevant to the task at hand, the simpler (and computationally faster) default algorithms can be used.

Integration of ACT-R and neural models of interval timing

The passing of time plays an important role in cognitive functioning, ranging from effects of impatience (e.g., [24]) to intricate temporal dependencies in the operation of complex machinery, and vice versa [25]. Until recently, computational models of behavior on cognitive tasks often ignored the effects of the passing of time on behavior, and computational models of interval timing ignored that interval timing is almost never relevant in isolation, but provides a supporting role for optimal performance. Recently, we proposed a first extension to ACT-R that could be used by cognitive models to keep track of subjective time [26–28]. This model, based on an information processing model of interval timing, implements a highly abstract conceptualization of the internal clock. Although this provides a useful approximation of human subjective timing, the high level of abstraction renders this model extremely flexible, making it almost impossible to falsify [29••]. To increase modeling rigor, we have recently started to replace the information-processing account with a neurobiological model of interval timing, the Striatal Beat Frequency (see for discussions [30,31]) model. This model reserves an important role for the striatum, an assumption shared widely although the functional interpretation sometimes differs (e.g., [32]). Interestingly, initial work has demonstrated that at a higher, function level the information processing and biologically plausible model come to similar predictions [29••]. However, the integration of SBF in a modeling framework such as ACT-R allows for testing much more detailed predictions, such as the role of memory consolidation on interval timing [33] or how interval timing can influence perceptual processes by means of predictive coding [34].

Conclusion

The question of the right level of abstraction has become more prominent since recent cognitive modeling efforts often aim to explain behavioral and neurophysiological data in a single framework ([35], e.g., [36••]). This trend entails that the data that a model should explain also yields different levels. Consequently, the need to include more (biological) detail in cognitive models becomes more evident ([37], see the special issue: [38], or [39••]).

The distinction that we draw between the model component of interest, and the remaining components raises a new question. That is, why is it necessary to model those aspects of behavior that one is not explicitly interested in the first place? There are three reasons for this. Firstly, it allows for generalization of mechanisms across domains. As Newell (1973) argued before us, behavior is the result of a single mind, and for this reason it is highly likely that different behaviors are the result of unique interplays of a limited set of mechanisms. Historically, this insight has led to the development of cognitive architectures such as ACT-R, but ultimately it could result in whole-brain models ([40••], e.g., [41,42•,43,44]).

Secondly, modeling non-focal components of a task is often useful to improve the model's prediction of those aspects of the data the cognitive model is focusing on. For example, the 'non-decision time' parameter in accumulator model is used to estimate the average duration of all non-decision processes. Obviously, a more detailed model can explain more variance related to the non-decision time, variance that without this extension could potentially be attributed, erroneously, to the decision process.

Thirdly, there could be a situation in which a researcher extends a pre-existing model to account for a novel finding. In this case it could be useful to maintain modeling aspects that are not the focus of the new model, for backwards compatibility, and future research in directions other than the current one.

Yet, working at different levels of abstraction does introduce new challenges. For example, although at higher levels of abstraction certain mechanisms are assumed to be completely independent, implementing these mechanisms at the neurobiological level might demonstrate that the mechanisms interact. For example, in the context of ACT-R, production rule execution and interval timing are assumed to be completely independent modules. However, more biologically plausible implementations of both modules suggest that the neural substrate driving these processes have a reasonable overlap. Although this violates the principles of modularity, it also provides new avenues for further study, as an updated architecture should acknowledge these interactions.

To conclude, although the first full-brain computational models have been proposed, the tasks to which they can be applied are still fairly limited, and the fit between model performance and empirical data is often more at a qualitative level. In the next decade, we expected therefore to see an increasing number of models in which different modeling abstractions are mixed to obtain the right balance between precision, neurobiological plausibility, and inspectability.

Conflict of interest

Nothing declared.

Acknowledgements

HvR was funded by EU FET grant (GA: 641100) TIMESTORM; JB by a NWO grant # 451-15-040, NT by ERC-StG grant 283597.

References and recommended reading

Papers of particular interest, published within the period of review, have been highlighted as:

- • of special interest
- •• of outstanding interest

1. Van Rija J, Van Rijn H, Hendriks P: **Cognitive architectures and language acquisition: a case study in pronoun comprehension**. *J Child Lang* 2010, **37**:731.

2. • O'Reilly R, Hazy T, Herd S: **The Leabra cognitive architecture: how to play 20 principles with nature and win!** In *The Oxford Handbook of Cognitive Science*. Edited by Chipman S., in press. Retrieved from http://psych.colorado.edu/~oreilly/papers/oreillyhazyherdip.pdf

This paper describes one of the first approaches to a whole-brain computational model based on sound neurobiological principles. It has also been used as a more detailed implementation of the visual module of the ACT-R cognitive framework, demonstrating the feasibility of the proposed approach.

3. • Engelmann F, Jäger LA, Vasishth S: *The Determinants of Retrieval Interference in Dependency Resolution: Review and Computational Modeling [Internet]*. 2016 In: *http://www.ling.uni-potsdam.de/~vasishth/pdfs/EngelmannJaegerVasishthR12016.pdf*.

This working paper describes in detail how extending an existing computational cognitive model with more detailed modules improves the explanatory power of the model. The computational models in this paper are available online and well documented.

4. Borst JP, Buwalda TA, van Rijn H, Taatgen NA: **Avoiding the problem state bottleneck by strategic use of the environment**. *Acta Psychol (Amst)* 2013, **144**:373-379.

5. van Maanen L, van Rijn H, Borst JP: **Stroop and picture–word interference are two sides of the same coin**. *Psychon Bull Rev* 2009, **16**:987-999.

6. Pitt M, Myung I: **When a good fit can be bad**. *Trends Cogn Sci* 2002, **6**:421-425.

7. Marr D, Poggio T: *From Understanding Computation to Understanding Neural Circuitry*. Massachutes Institute of Technology; AI Memo; 1976:: 357.

8. Newell A: *Unified Theories of Cognition*. Harvard University Press; 1990.

9. Anderson JR: **Spanning seven orders of magnitude: a challenge for cognitive modeling**. *Cogn Sci* 2002, **26**:85-112.

10. Gray WD (Ed): *Integrated Models of Cognitive Systems*. Oxford University Press; 2007.

11. Meeter M, Jehee J, Murre J: **Neural models that convince: model hierarchies and other strategies to bridge the gap between behavior and the brain**. *Philos Psychol* 2007, **20**:749-772.

12. Anderson JR: *How Can the Human Mind Occur in the Physical Universe?* Oxford University Press; 2007.

13. •• Anderson JR, Zhang Q, Borst JP, Walsh MM: **The discovery of processing stages: extension of Sternberg's method**. *Psychol Rev* 2016 http://dx.doi.org/10.1037/rev0000030.

This paper discusses how EEG data can be used to identify processing stages and procedures, providing additional information to modelers about when more detailed modules have to be considered.

14. Borst JP, Taatgen NA, van Rijn H: **What makes interruptions disruptive?** *Proc. 33rd Annu. ACM Conf. Hum. Factors Comput. Syst. - CHI '15*. 2015 http://dx.doi.org/10.1145/2702123.2702156.

15. Borst JP, Anderson JR: **The discovery of processing stages: analyzing EEG data with hidden semi-Markov models**. *Neuroimage* 2015, **108**:60-73.

16. Borst JP, Anderson JR. **A step-by-step tutorial on using the cognitive architecture ACT-R in combination with fMRI data**. *J Math Psychol*, in press http://dx.doi.org/10.1016/j.jmp.2016.05.005.

17. Borst JP, Nijboer M, Taatgen NA, van Rijn H, Anderson JR: **Using data-driven model-brain mappings to constrain formal models of cognition**. *PLoS One* 2015, **10**:e0119673.

18. Nijboer M, Borst JP, van Rijn H, Taatgen NA: **Single-task fMRI overlap predicts concurrent multitasking interference**. *Neuroimage* 2014, **100**:60-74.

19. Borst JP, Taatgen NA, van Rijn H: **The problem state: a cognitive bottleneck in multitasking**. *J Exp Psychol Learn Mem Cogn* 2010, **36**:363-382.

20. Nijboer M, Borst JP, van Rijn H, Taatgen N: **Contrasting single and multi-component working-memory systems in dual tasking**. *Cogn Psychol* 2016, **86**:1-26.

21. van Maanen L, van Rijn H, Taatgen NA: **RACE/A: an architectural account of the interactions between learning, task control, and retrieval dynamics**. *Cogn Sci* 2012, **36**:62-101.

22. Usher M, McClelland J: **The time course of perceptual choice: the leaky, competing accumulator model**. *Psychol Rev* 2001, **108**:550-592.

23. van Maanen L, van Rijn H: **Accounting for subliminal priming in ACT-R**. *8th Int. Conf. Cogn. Model.*. 2007. [no volume].

24. Boehm U, Hawkins GE, Brown S, van Rijn H, Wagenmakers E-J: **Of monkeys and men: impatience in perceptual decision-making**. *Psychon Bull Rev* 2015 http://dx.doi.org/10.3758/s13423-015-0958-5.

25. van Rijn H: **It's time to take the psychology of biological time into account: speed of driving affects a trip's subjective duration**. *Front Psychol* 2014, **5**:1028.

26. Taatgen NA, van Rijn H, Anderson JR: **An integrated theory of prospective time interval estimation: the role of cognition, attention, and learning**. *Psychol Rev* 2007, **114**:577-598.

27. Taatgen NA, van Rijn H: **Traces of times past: representations of temporal intervals in memory**. *Mem Cognit* 2011, **39**:1546-1560.

28. van Rijn H, Taatgen NA: **Timing of multiple overlapping intervals: how many clocks do we have?** *Acta Psychol (Amst)* 2008, **129**:365-375.

29. •• van Rijn H, Gu B, Meck WH: **Dedicated clock/timing-circuit theories of time perception and timed performance**. *Adv Exp Med Biol* 2014, **829**:75-99.

This paper describes how and cognitive modeling framework and an information processing model of a specific phenomenon can be combined, and how replacing the information processing model with a biologically plausible model increases the model's validity.

30. Lusk NA, Petter EA, MacDonald CJ, Meck WH: **Cerebellar, hippocampal, and striatal time cells**. *Curr Opin Behav Sci* 2016, **8**:186-192.

31. Merchant H, Harrington DL, Meck WH: **Neural basis of the perception and estimation of time**. *Annu Rev Neurosci* 2013, **36**:313-336.

32. Mello GBM, Soares S, Paton JJ: **A scalable population code for time in the striatum**. *Curr Biol* 2015, **25**:1113-1122.

33. Gu B-M, van Rijn H, Meck WH: **Oscillatory multiplexing of neural population codes for interval timing and working memory**. *Neurosci Biobehav Rev* 2015, **48**:160-185.

34. van Rijn H: **Accounting for memory mechanisms in interval timing: a review**. *Curr Opin Behav Sci* 2016, **8**:245-249.

35. Matthews WJ, Terhune DB, van Rijn H, Eagleman DM, Sommer MA, Meck WH: **Subjective duration as a signature of coding efficiency: emerging links among stimulus repetition, predictive coding, and cortical GABA Levels**. *Timing Time Percept Rev* 2014:1.

36. •• Turner B, Forstmann BU, Love BC, Palmeri J, van Maanen L: **Approaches to analysis in model-based cognitive neuroscience**. *J Math Psych*, in press.

This paper discusses how computational models derived from cognitive neuroscience data can be validated by linking neural and behavioral data.

37. de Hollander G, Forstmann B, Brown S: **Different ways of linking behavioral and neural data via computational cognitive models**. *Biol. Psychiatry: Cogn. Neurosci. Neuroimaging* 2016, **1**:101-109.

38. Purcell BA, Heitz RP, Cohen JY, Schall JD, Logan GD, Palmeri TJ: **Neurally constrained modeling of perceptual decision making**. *Psychol Rev* 2010, **117**:1113-1143.

39. •• Dehaene S, Dudai Y, Konen C: **Cognitive architectures**. *Neuron* 2015, **88**:1.

This is the editorial of an excellent special issue of Neuron on Cognitive Architectures. Most of the cognitive models presented in this special issue model all aspects of the cognitive task at a single level, making them potentially useful to be integrated in a cognitive framework consisting of different levels of abstraction.

40. •• Dehaene S, Meyniel F, Wacongne C, Wang L, Pallier C: **The neural representation of sequences: from transition probabilities to algebraic patterns and linguistic trees**. *Neuron* 2015, **88**:2-19.

This paper describes how one might explain how the lowest levels of abstraction can generate relatively complex forms of cognitive performance for which it is typically claimed that a more abstract processing system is required.

41. Eliasmith C, Stewart TC, Choo X, Bekolay T, DeWolf T, Tang Y, Tang C, Rasmussen D: **A large-scale model of the functioning brain**. *Science* 2012, **338**:1202-1205.

42. • Eliasmith C: *How to Build a Brain: A Neural Architecture for Biological Cognition*. 2013.

This book, and the matching paper [41], describe a whole brain model that has proven relatively succesful in explained a large range of phenomena although for some phenomena a more abstract level of computational modeling is still necessary.

43. Shanahan M: **A cognitive architecture that combines internal simulation with a global workspace**. *Conscious Cogn* 2006, **15**:433-449.

44. Zylberberg A, Dehaene S, Roelfsema PR, Sigman M: **The human turing machine: a neural framework for mental programs**. *Trends Cogn Sci* 2011, **15**:293-300.