ELSEVIER

Journal of Cognitive Systems Research 1 (2001) 000–000

Cognitive Systems RESEARCH

www.elsevier.com/locate/cogsys

Modeling selective attention: not just another model of Stroop (NJAMOS)

Marsha C. Lovett*

Department of Psychology, Carnegie Mellon University, Pittsburgh, PA 15213, USA

Received 1 March 2001; accepted 1 September 2001

Abstract

The Stroop effect has been studied for more than sixty years, and yet it still defies a complete theoretical account. The model presented here offers a new approach that integrates several explanations of the Stroop phenomenon into a hybrid model. Because this model is built within the ACT-R cognitive architecture (Anderson & Lebiere, 1998), it applies a generic, pre-specified set of mechanisms for learning and performance to the particulars of the Stroop task. Besides fitting a variety of already published experimental results, the model offers the potential to capture strategic variation in what is typically considered a low-level attentional phenomenon. © 2001 Published by Elsevier Science B.V.

Keywords: Cognitive architectures; Selective attention; Stroop task; Hybrid models

1. Modeling selective attention: not just another model of Stroop (NJAMOS)

The Stroop effect has been studied for more than sixty years (Stroop, 1935), and yet it still defies a complete theoretical account. One explanation for the apparent lack of progress is that so much empirical research has been conducted using this basic paradigm that what we now call the 'Stroop effect' is actually a compendium of results derived from a multitude of manipulations applied to a family of Stroop-like tasks! The current article focuses on a select set of Stroop results in order to introduce the model NJAMOS. NJAMOS offers a new theoretical account that integrates several explanations of the Stroop phenomenon into a hybrid model. Specifically, NJAMOS performs competitive, parallel retrieval of information within a goal-based, sequential cognitive processor. NJAMOS is built within the ACT-R cognitive architecture (Anderson & Lebiere, 1998), so it applies a general, pre-specified set of learning and performance mechanisms to the particulars of the Stroop paradigm. Moreover, NJAMOS is unique among models of ('low level') attentional phenomena in that it allows for ('high level') strategic variability.

The organization of the paper is as follows. First a description of the Stroop phenomenon is presented. Then, major theoretical features of other computational models are reviewed. Next, the NJAMOS model is described and fit to a selection of relevant

*Tel.: +1-412-268-3499; fax: +1-412-268-2798.
E-mail address: lovett@cmu.edu (M.C. Lovett).

1389-0417/01/$ – see front matter © 2001 Published by Elsevier Science B.V.
PII: S1389-0417(01)00045-6

data. Finally, the potential of NJAMOS for accommodating other results is discussed along with a more general view on evaluating Stroop models.

1.1. The basic phenomenon

The Stroop effect offers a window onto the processes of selective attention in that stimuli with two prominent dimensions are presented in a task where one dimension must be processed and the other ignored. Typically, the stimuli are words, and the two dimensions are the form of the word and the color of the ink in which it is written. The task, then, is either to name the ink color or to read the word. The basic Stroop manipulation varies the relationship between the meaning of the word and the color of the ink to be congruent (e.g., the word ‘red’ printed in red ink), conflicting (e.g., the word ‘blue’ printed in red ink), or neutral (e.g., the word ‘dog’ or a string of ‘X’s printed in red ink). A robust result emerges: for color naming, there is interference in the conflicting case and (usually) facilitation in the congruent case, but for word reading, there is no (or very little) effect of the congruency of this relationship.

Fig. 1 shows a typical data set (along with the NJAMOS predictions to be discussed later). The interference and facilitation in color naming can be seen by the shifts in the color-naming curve as a function of congruency. The lack of such effects for word reading are shown by the relatively flat line for this condition. These results suggest an asymmetry in selective attention, namely, that participants are strongly influenced by the word when naming the ink color but that they can ignore the ink color when reading.

2. Theoretical accounts of the Stroop effect

Two different views of the Stroop effect cover much of the theoretical work in this area. The ‘horse-race’ view highlights the overall difference in speed of processing for words versus colors (see Fig. 1, separation of the two curves) and assumes a response bottleneck. This view implies that the pattern of interference depends on the relative arrival of word versus color information to the response stage: whichever kind of information arrives first will produce interference for the other. Because word reading is, on average, faster than color naming, this view predicts the asymmetry of words interfering with colors but not vice versa.

The other view of the Stroop effect highlights the different levels of automaticity people have acquired for processing the two stimulus dimensions. Because word reading is so highly practiced among typical Stroop experiment participants, it is more automatic than color naming. This greater automaticity implies that reading requires fewer attentional resources and hence interferes more easily with color naming.

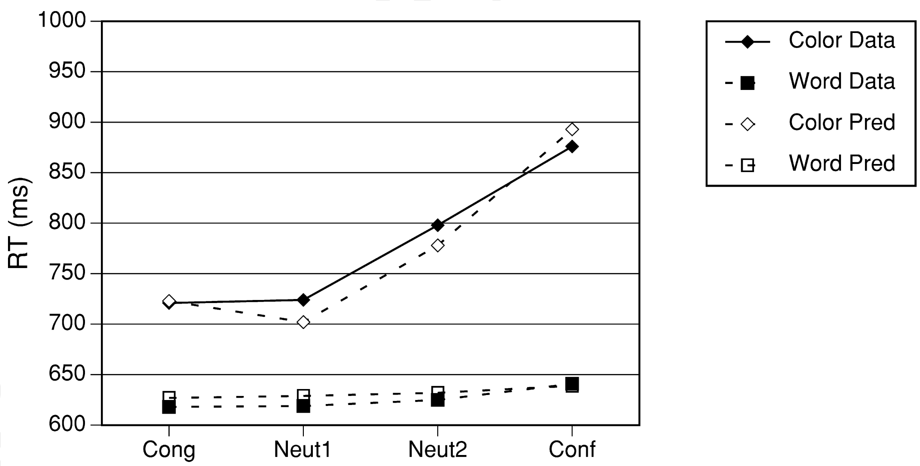

Fig. 1. Reaction times for standard Stroop experiment. ‘Neut1’ refers to a string of ‘X’s in colored ink for color naming and a word printed in black ink for word reading; ‘neut2’ refers to a non-color word in colored ink for color naming.

The key similarity between the two views is that they both emphasize parallel processing of the two stimulus dimensions. Not surprisingly, then, the dominant computational accounts of Stroop phenomena have been implemented within connectionist models. The key difference between the two views is whether relative speed or automaticity is considered the main determiner of interference effects. Note that manipulations of stimulus-onset-asynchrony (SOA) alter the timing of processing, predicted to be important under the horse-race view, and manipulations of practice can alter the automaticity of processing that is predicted to be important under the automaticity view. Thus, experiments that employ these manipulations are crucial for evaluating specific theoretical accounts of Stroop phenomena and hence for evaluating particular computational models as well.

2.1. Existing computational models

Three existing computational models will be discussed. Two of these are connectionist models (Cohen, Dunbar & McClelland, 1990; Phaf, van der Heijden & Hudson, 1990), and the other is built as a production system (Roelofs, 2000). The Cohen et al. (1990) model was designed to capture an automaticity account of Stroop. It has input nodes that connect to hidden-layer nodes which connect to response nodes via either a word-reading or color-naming pathway. Task-control nodes are connected to both pathways to 'gate' the information processing in a way that biases towards the instructed task. To represent the greater automaticity of reading, stronger weights are given along the word-reading pathway, making reading the presented word less sensitive to the congruency relationship of the simultaneously presented color than color-naming is to the congruency relationship of the simultaneously presented word. The Phaf et al. (1990) model was built as the Stroop-task extension to a general model of selective visual attention. This model's network architecture differs for word reading and color naming in that there are direct input–output connections for word reading but not for color naming. This difference removes the intermediate (hidden layer) processing for word reading which reduces the potential for semantically based interference in that task.

The Roelofs (2000) model of Stroop phenomena is built as an extension to the WEAVER++ model of word production (Levelt, Roelofs & Meyer, 1999). As such, this Stroop model specifies separate processing stages for lemma retrieval, word-form encoding, etc. Like the Phaf et al. (1990) model, it builds in a word-reading advantage by requiring fewer processing steps for that task, hence enabling different-sized congruency effects between word reading and color naming.

2.2. These models' fits to data

These three models all account for the basic results in Fig. 1, but how do they fare in predicting other important Stroop results? As mentioned above, one critical Stroop manipulation varies SOA. Glaser & Glaser's (1982) Experiment 1 did this by presenting the two stimulus dimensions spatially separated so that the onset of word and color information could be lagged. SOAs were varied from −400 ms to +400 ms, where negative SOA means preexposure of the irrelevant information. Fig. 2a presents these data.

All three models have simulated these data to varying degrees of success. They all predict the standard Stroop result at 0 SOA, and they all show a reduced Stroop effect for color naming at positive SOAs, a result that is consistent with the observed data. There are several areas of misfit, however. Both connectionist models show a small but notable congruency effect for word reading at negative SOAs, even though this is not present in the data. In addition, the Cohen et al. model predicts that, in color naming, the congruency effect will be largest at negative SOAs, when it is actually reduced here relative to 0 SOA. The Roelofs model has neither of these problems but fails to fully capture the smoothness of the increase in interference for color naming from −400 to 0 SOA.

The other critical manipulation mentioned above is degree of practice. MacLeod and Dunbar (1988) devised a clever variant of the Stroop task in which they presented stimuli with the key dimensions of shape and color and asked participants to either name the shape or the color. So that congruency of the

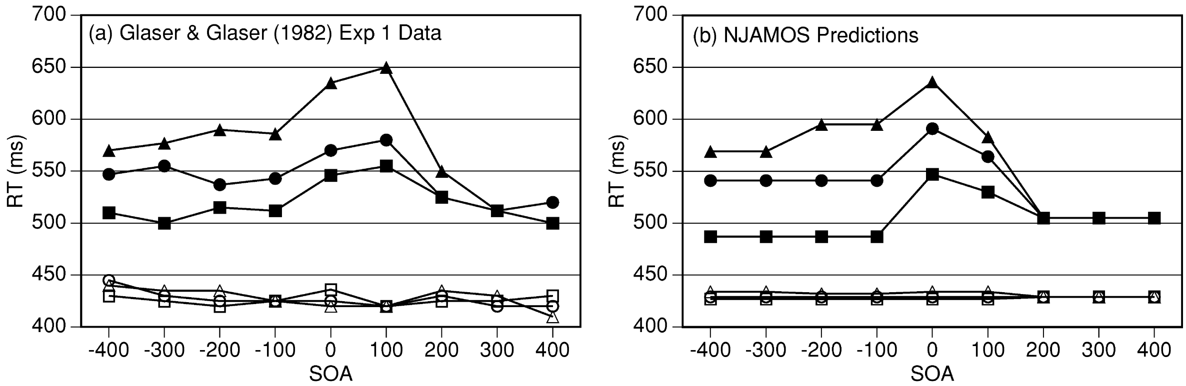

Fig. 2. Data from Glaser and Glaser (1982) (a) and NJAMOS predictions (b). RTs plotted for color naming (solid symbols) versus word reading (open symbols) and for conflicting (triangle), neutral (circle), and congruent (square) stimuli.

shape-color relationship could be varied, the shapes came from a fixed set of unfamiliar shapes, and participants were trained to name each shape with a specific color word (e.g., the irregular hexagon is named ‘red’). Manipulating practice involved giving participants 20 sessions of shape training and measuring their Stroop performance along the way. Fig. 3a presents these data.

The Cohen et al. (1990) model nicely captures the early (training session 1) pattern of colors interfering with shape naming and not vice versa as well as the late (session 20) pattern of shapes interfering with color naming and hardly vice versa. This reversal is accomplished by the model’s weight-learning mechanism that gradually strengthens the shape-naming pathway with practice. An interesting transition point in the data (session 5) shows considerable interference for both tasks. However, the Cohen et al. model makes its transition by predicting no interference in either task at this point.

Regarding practice effects, the Phaf et al. and Roelofs models are essentially silent. In both, the word-reading advantage is implemented as a qualitative change (a ‘short cut’ for word-reading), and there is no specific learning mechanism presented.

2.3. Summary of existing models

To be competitive, any computational model of Stroop effects must demonstrate some added value. What do each of these models offer? The Cohen et al. model offers the best coverage of Stroop-related

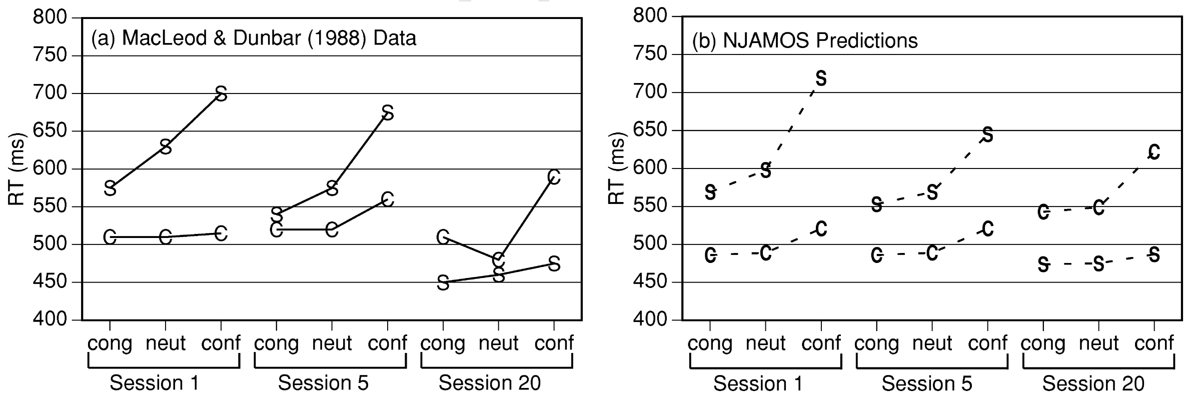

Fig. 3. Data from MacLeod and Dunbar (1988) (a) and NJAMOS predictions (b). In each panel, RTs for shape naming (S) and color naming (C) are plotted separately across congruency and training manipulations.

data at present and does so with very simple, natural connectionist mechanisms for strengthening connections and gating activation along the two pathways. This trades off against the model's acknowledged exclusion of other processes (e.g., habituation, strategy choice, etc.) that could help address its known areas of misfit. Note that, while the model was designed to account for Stroop phenomena per se, its basic form and mechanisms have been applied to a slightly broader set of paradigms (Cohen & Huston, 1994).

The Phaf et al. model and the Roelofs model were both devised to account for Stroop phenomena within a more general model of a related task. These models demonstrate how processes that play a central role in another task can help explain certain Stroop effects. However, this breadth of application tends to dilute the models' coverage of Stroop results themselves, and neither model includes a learning mechanism that could explain its shortcut pathways for reading.

3. A hybrid model, NJAMOS

There are several sources of added value in NJAMOS. First, it is built within the ACT-R cognitive architecture (Anderson & Lebiere, 1998), so it applies a general, pre-specified set of learning and performance mechanisms to the particulars of the Stroop paradigm. As such, it has many of the benefits of a general approach in that the mechanisms that drive its predictions have been shown to produce accurate predictions for a variety of other tasks. This general-mechanisms approach also offers guidance and constraint in developing the specific task model beyond what the empirical data on that task can provide, i.e., the model must be built to fit the relevant data within a given, pre-specified structure (cf. Newell, 1990).

Second, building a specific task model within a cognitive architecture such as ACT-R still allows (and arguably facilitates) a focus on capturing as many experimental results associated with that particular task as possible. Indeed, with the architecture taken as given, model development involves specifying the knowledge used in performing the current task, some of which may come from participants' prior experience and some of which may come from their exposure to the task itself. Third, NJAMOS is unique among models of ('low level') attentional phenomena in that it allows for ('high level') strategic variability. Evidence suggests that participants can (and do) apply different strategies when performing the Stroop task (e.g., Chen & Johnson, 1991; Logan, Zbrodoff & Williamson, 1984, Lovett, 2001). This strategic variation may be considered a source of qualitative differences in performance from trial to trial that this model (unlike many others) can produce. Note that NJAMOS can also produce quantitative differences in performance from run to run (i.e., simulated participant to simulated participant) by varying the ACT-R parameter associated with individual differences in working memory capacity (Lovett, Daily & Reder, 2000). The latter effect maps onto individual differences found in the Stroop task and will be discussed below.

3.1. NJAMOS model overview

NJAMOS specifies the knowledge relevant for performing Stroop-like tasks. In ACT-R, this is comprised of a set of production rules and a network of declarative chunks. Each production rule takes the form 'IF ⟨conditions⟩ THEN ⟨actions⟩' and has an associated measure of utility. A production rule's utility is learned through past experiences; it provides an index of how effective the production rule has been. In ACT-R, when more than one production rule's conditions are met, these rules compete on the basis of their utility values. The production rule with the highest utility (after some noise is added) is applied, and its actions are executed. In practice, this means that the system can learn to choose (and prefer) more effective strategies.

A declarative chunk is represented as a template with slots for related information. At the sub-symbolic level, each chunk has a base-level activation representing its overall accessibility. A network of chunks is specified by directed, pairwise connections, each with a strength of association. These strengths influence how strongly one chunk (in the current focus of attention) cues the availability of another chunk. The ACT-R equation for total activation of chunk i is

$$A_i = B_i + SW_jS_{ji}$$

where B_i is the base-level activation of chunk i, W_j is the amount of attention focused on chunk j, and S_{ji} is the strength of association from chunk j to i. This quantity then translates into the chunk's likelihood of being retrieved and latency to retrieval. The latency measure is more relevant in NJAMOS; it follows the following function

$$L_i = F\ \mathrm{e}^{-Ai}$$

where F is a latency-scale factor.

Some specific examples of the knowledge units specified in NJAMOS (and presented in English for readability) are presented in Fig. 4. The first two examples in Fig. 4 are production rules that would initiate processing of the color and word dimensions of the stimulus, respectively. The third example is a declarative chunk representing information about the color black. The last example is a specification of a goal chunk for the color-naming task after the stimulus has been perceived but before any color concept has been retrieved (either by processing the word or the color). In ACT-R, processing is in large part driven by the current goal. The goal is represented like other declarative chunks but it is considered to be the current focus of attention. A fixed, limited amount of source activation (W in ACT-R) is shared among the different slots of the current goal (W_j for each slot). This share of source activation is spread to connected chunks (proportionally to the strength of the connection) and added to the receiving chunk's base-level activation to comprise that chunk's total activation.

Given this specification of knowledge and ACT-R's fixed mechanisms, two important features of NJAMOS's performance follow. First, production rule choice favors word reading (even when the instructed task is color naming). This is because the word-processing production (see Fig. 4) is general enough that it applies whenever there is a word-like stimulus and because its utility value is taken to be much higher than that for color naming. The generality of the word-reading production rule represents the extensive practice people have with reading that leads words to be processed regardless of the current goal; this kind of general rule could be learned under ACT-R's production-learning mechanism. The higher utility of the word-reading production rule represents the fact that reading often helps people to achieve their goals; ACT-R's utility-learning mechanism would naturally produce this higher value given extensive practice.[1]) Production-rule choice provides NJAMOS with a degree of sequentiality as each production rule in this model

[1] Note that the model's performance does not require that color naming be constrained by a specific color-task goal. As long as word reading has a higher utility, both word-reading and color-naming productions could fire under a generic goal.

IF the goal involves a color-naming task, and the stimulus color has been encoded,
THEN retrieve the associated color concept.

IF the goal involves processing a stimulus [generic goal] and the stimulus has word-like qualities, THEN retrieve the associated word concept.

Color-association-chunk1
Stimulus: 
Feature: color
Concept: black

Current-goal
Stimulus-word: "green"
Stimulus-color: 
Task: color

Fig. 4. Example production rules, chunks, and goal used in NJAMOS (presented in pseudo-code for readability).

tends to focus on a single dimension (word or color). This also provides NJAMOS with strategic variability: processing of the word may or may not precede processing of the color under the color-naming task. Either way, a 'check' production works to verify that the correct dimension has been processed before enabling a response.

Parallel processing also plays a role in producing the second important feature of NJAMOS performance, namely, that both dimensions of the stimulus influence the retrieval of declarative chunks. This is because all of the slots in the current goal provide contextual cues to influence the total activation of to-be-retrieved chunks (see total activation equation). Some of these slots represent the stimulus features (word and color); others may represent a concept retrieved during previous processing (e.g., if the word were partially processed because the word-reading production had won over color naming, then the concept of the word would provide another piece of context present in the focus of attention during retrieval of the ink color). Recall that these contextual effects are proportional to the strengths of association between the relevant chunks. In NJAMOS, pairs of chunks that involve the same color concept (e.g., a color-association chunk for blue and a word-association chunk for blue) are given positive strengths of association, and chunks that involve different color concepts have negative strengths of association. These positive and negative strengths contribute to the facilitation and interference effects. NJAMOS produces larger facilitation and interference effects in color naming (see Fig. 1) predominantly because (a) color-association chunks are given a lower base-level activation which, in the nonlinear latency equation (see above), makes them not only slower overall but more sensitive to contextual cueing and (b) the concept of the word will often have been retrieved into the goal before retrieving the color (but not vice versa), magnifying the contextual effect of the word dimension of the stimulus.

Fig. 5 presents a sketch of the flow of control in NJAMOS, summarizing these relationships and indicating other examples of production-rule choice and parallel chunk retrievals. Note that any branching in this diagram represents an opportunity for trial-to-trial strategic variability.

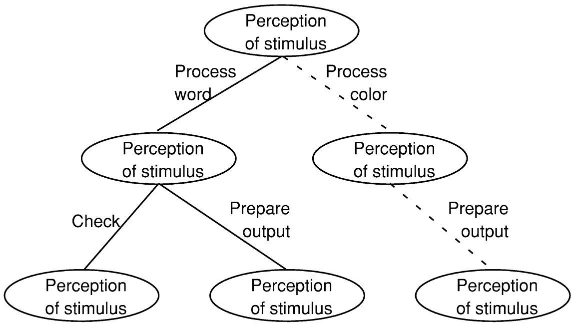

Fig. 5. Flow of control in NJAMOS for color naming with standard stimuli. Lines represent production rules (note choice points), and ovals represent parallel retrievals that modify the current focus of attention. From the bottom left oval (*), flow continues along the dashed line path.

3.2. Fitting NJAMOS to the data

This section describes how NJAMOS was fit to five separate experiments comprising a total of 92 data points (including Fig. 1). For ease in estimating a set of best-fitting parameters, a mathematical description of NJAMOS was used. The same parameter settings were used to predict performance across all five experiments[2] with the following exceptions. For each experiment two parameters were varied. One of these is the latency-scale factor F (see latency equation above), and the other was a latency intercept. These parameters appeared necessary to capture the experiment-to-experiment variation in latencies that appears even under very similar manipulations and experimental designs (e.g., compare RTs in Figs. 1 and 6). In addition, in two cases, another single parameter was varied when needed for a particular experiment's manipulation; those cases will be discussed below. To preview the quality of the NJAMOS fit, all 92 data points were fit by varying a

[2]Parameters taken as fixed (i.e., not optimized to fit the data) but whose values departed from ACT-R's default values include base-level activation for word-association chunks (set at 2) and for color-association chunks (set at 0), S_{ji}'s between same-color chunks (set at 1.5), between different-color chunks (set at −1.5), between same-task chunks (set at 0.6), and between different-task chunks (set at 0).

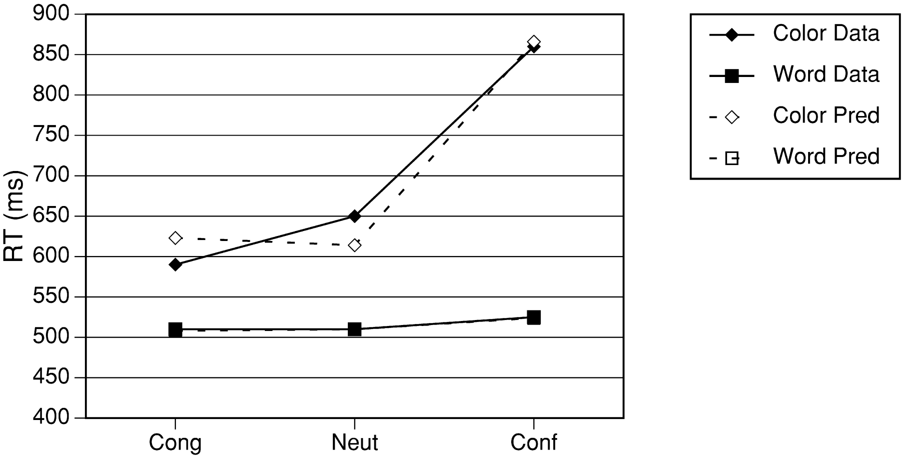

Fig. 6. Dunbar and MacLeod (1984) data and model fit.

total of 12 parameters for an R^2 of 0.95 and MSE of 1114. In essence, the goal was to test whether NJAMOS can produce reasonable accounts of these experiments without much parameter varying.

3.3. Dunbar and MacLeod (1984)

Fig. 6 shows the data and NJAMOS predictions for Dunbar and MacLeod's experiment (1984, p. 62). The design is similar to that presented in Fig. 1, with the neutral condition corresponding to 'Neut1' a string of X's for the color-naming task. Both the data and model show the standard Stroop results. As discussed above, NJAMOS shows little facilitation or interference in word reading because the parallel retrieval process in this task is based on retrieval of highly practiced word-reading chunks whose base-level activations are little affected by added contextual activation. In contrast, the lower base-level activations of color-naming chunks are significantly impacted by spreading activation (positive/facilitatory or negative/inhibitory) from elements currently in the focus of attention (e.g., the word stimulus or even the word-concept already retrieved). This differential effect of added contextual activation on word-reading chunks versus color-naming chunks highlights the nonlinear relationship between activation and retrieval latency which is pre-specified in ACT-R (see latency equation above).

3.4. Glaser and Glaser (1982)

Fig. 2b shows the predictions of NJAMOS for the Glaser and Glaser (1982) experiment described above. Unlike previous models, NJAMOS appears to capture (at least qualitatively) three effects in the data: (1) no congruency effect for word reading at any SOA, (2) increase in overall latency for color naming as SOA goes from −400 to 0, and (3) increase and then decrease in interference effects from −400 to 0 to +400 SOA. Much of this match to the data is driven by the model's tendency to choose word reading first, even on color naming trials. Note, however, that the size of the interference effect is not monotonically increasing from −400 to 0 as it is in the data and the facilitation effect in color naming is overpredicted. Under the parameter-fitting constraints applied here (2 free parameters for this set of data), this is not too surprising. The fit can be improved by allowing more parameters to vary (e.g., informal sensitivity analyses revealed that modifying one or two S_{ji} values has an impact).

3.5. MacLeod and Dunbar (1988)

Fig. 3b presents the NJAMOS predictions for MacLeod and Dunbar's (1988) experiment. Note that this is a case where an additional parameter was required by the design of the experiment, namely, the initial base-level activation for the shape-name-association chunks was fit for this experiment alone (because no other experiments employed the novel shape-naming task). From this initial value, ACT-R's learning mechanism naturally increased this base-level activation as the model received additional practice by retrieving the shape-name-association chunks during training. Note that these increases in base-level activation make the shape chunks less and less susceptible to context effects (just as the model fit above demonstrated with high base-level activations for word-association chunks). Also, NJAMOS can learn to adjust its strategy of trying to name the color first (regardless of task) as the two competing (shape and color-processing) productions adjust their utilities. Indeed, over time, the model learns to process shape information first, rather than color information. This increases the proportion of trials in which shape information is present in the goal and can serve as a source of contextual activation impacting (usually inhibiting in the case of conflict stimuli) the task of color naming.

NJAMOS shows three main results in this experiment: (1) at session 1, there is more interference and facilitation in shape naming than in color naming, (2) at session 5, there is some facilitation and interference in both tasks, and (3) at session 20, there is more facilitation and interference in color naming than shape naming. Note that the Cohen et al. model captured only the first and third of these. Nevertheless, NJAMOS appears to have difficulty capturing the increase in the interference effect for color naming between sessions 1 and 5. This is probably caused by the model's overprediction of interference for color naming in session 1 (30 ms in the predictions versus 5 ms in the data). It is possible that adjusting the production utility parameters would address this issue, but as mentioned above, the motivation for the current model fits was to demonstrate that NJAMOS could produce reasonable model fits with a relatively small number of free parameters.

3.6. *Cohen et al. (1992)*

Table 1 presents Stroop data from schizophrenic patients and matched controls along with the NJAMOS predictions. The main result is a larger interference effect (for color naming) among the

Table 1
Data and model fit to schizophrenic data summarized in Cohen et al. (1992)

	Neut.-word	Neut.-color	Conf.-color
Patient data	530	797	1467
Control data	420	603	1037
Patient pred.	456	764	1345
Control pred.	440	661	1137

patients. While schizophrenia is a very complex condition, it is associated with a deficit in working memory. To capture this NJAMOS simply uses an adjusted *W* parameter. This was the extra parameter varied to fit these data. As the table shows, reducing *W* produces the effect as well as raising RTs among the patients. Notice that in the total activation equation above, a reduced *W* will lower the activation of target chunks (relative to control participants with $W = 1$). For word reading, this does not have much of an effect on the latency measure because of its nonlinearity, but for color naming, this difference between controls and patients is accentuated by the nonlinearity of the retrieval latency function.

4. Conclusions

The above model fits demonstrate that NJAMOS can fit a variety of Stroop results, when constrained by the ACT-R architecture and a limited number of free parameters. In many other, less studied paradigms, such a demonstration of a single model's fit to five separate experiments might offer an adequate account of the phenomenon. In the case of Stroop, however, there are many more results associated with the paradigm. MacLeod's (1991) review of the Stroop literature lists 18 key results that may serve as a core, but even that is a distillation of the many, varied manipulations — and corresponding results — this paradigm has wrought. This leads to a tension in model development and evaluation: How to balance a model's coverage of the broad span of Stroop results with its parsimony? An added dimension of evaluation that cognitive architectures have added to the mix is whether the model itself is a part of a broader system that can account for more than behavior in Stroop tasks?

NJAMOS aspires for all of these desiderata:

coverage, parsimony, and generality. At this point, its coverage is not yet large, but it is growing. For example, semantic gradient effects (e.g., Dalrymple-Alford, 1972) and trial-type frequency effects (e.g., Tzelgov, Henik & Berger, 1992) can be captured under its existing mechanisms. Its parsimony may be measured by the relatively small proportion of free parameters or by the small set of productions and chunks required to specify the model's incoming knowledge. Its generality is represented by the many related models that have been built within the ACT-R architecture using only different starting knowledge sets (where different knowledge is required for performing different experimental tasks). That is, the learning and performance mechanisms used by NJAMOS were not built for this Stroop model alone but rather have been used and tested in a wide range of models covering a broad spectrum of cognitive tasks.

That said, it is still quite difficult to generate adequate quantitative measures of coverage, parsimony, and generality that can be used to make across-model comparisons. For example, not all models will be set the task of fitting the same body of results, and in the case of Stroop, there are many different subsets of this large literature. One reasonable approach exemplified in MacLeod's review is to use a fixed set of results (i.e., the eighteen listed in that paper's appendix) for a standardized comparison of current models. But when the distinctions between model fits come to quantitative, not just qualitative, differences, this requires that all the modelers take the same fixed set of results to heart and do the work of fitting. Unfortunately, even in that best case scenario, comparisons will be complicated by the question of whether all results should be treated as equally important. And, then, there is the issue of adjusting a model's goodness of fit for its level of complexity. Here, the dimension of parsimony must come into play, but there is no gold standard for measurement, and the process of comparing model complexity is almost impossible in the case of models built within different frameworks (e.g., connectionist versus hybrid). Instead of giving up in light of these challenges, the best course may be to use the vast Stroop literature as an impetus for aiming to understand not only how well a given model can capture a given result but why.

Acknowledgements

I thank the International Conference on Computational Modeling reviewers who provided helpful comments on a previous version.

References

Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*, Erlbaum, Mahwah, NJ.

Chen, J., & Johnson, M. (1991). The Stroop congruency effect is more observable under a speed strategy than an accuracy strategy. *Perceptual & Motor Skills 73*, 67–76.

Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: a parallel distributed processing account of the Stroop effect. *Psychological Review 97*, 332–361.

Cohen, J. D., & Huston, T. A. (1994). Progress in the use of interactive models for understanding attention and performance. In: Umilta, C., & Moscovitch, M. (Eds.), Attention and performance 15: conscious and nonconscious information processing, MIT Press, Cambridge, MA, pp. 453–476.

Cohen, J. D., Servan-Schreiber, D., & McClelland, J. (1992). A parallel distributed processing approach to automaticity. *American Journal of Psychology*, 239–269.

Dalrymple-Alford, E. C. (1972). Associative facilitation and interference in the Stroop color–word task. *Perception and Psychophysics 11*, 274–276.

Dunbar, K., & MacLeod, C. M. (1984). A horse race of a different color: Stroop interference patterns with transformed words. *Journal of Experimental Psychology: Human Perception and Performance 10*, 622–639.

Glaser, M. O., & Glaser, W. R. (1982). Time course analysis of the Stroop phenomenon. *Journal of Experimental Psychology: Human Perception and Performance 8*, 875–894.

Levelt, W. J. M., Roelofs, A., & Meyer, A. S. (1999). A theory of lexical access in speech production. *Behavioral & Brain Science 22*, 1–75.

Logan, G. D., Zbrodoff, N. J., & Williamson, J. (1984). Strategies in the color–word Stroop task. *Bulletin of the Psychonomic Society 22*, 135–138.

Lovett, M. C., Daily, L. Z., & Reder, L. M. (2000). A source activation theory of working memory. *Cognitive Systems Research 1*, 99–118.

MacLeod, C. M. (1991). Half a century of research on the Stroop effect: an integrative review. *Psychological Bulletin 109*, 163–203.

MacLeod, C. M., & Dunbar, K. (1988). Training and Stroop-like interference: evidence for a continuum of automaticity. *Journal of Experimental Psychology: Learning, Memory, and Cognition 14*, 126–135.

Newell, A. (1990). *Unified theories of cognition*, Harvard University Press, Cambridge, MA.

Phaf, R. H., van der Heijden, A. H. C., & Hudson, P. T. W. (1990).

SLAM: a connectionist model for attention in visual selection tasks. *Cognitive Psychology 22*, 273–341.

Roelofs, A. (2000). Control of language: a computational account of the Stroop asymmetry. In: Proceedings of the Third International Conference on Cognitive Modeling, Universal Press, Veenendaal, The Netherlands.

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology 18*, 643–662.

Tzelgov, J., Henik, A., & Berger, J. (1992). Controlling Stroop effects by manipulating expectations for color words. *Memory and Cognition 20*, 727–735.