Inhibition of Task Sets

Ulrich Mayr

University of Oregon

Does our cognitive system use inhibition to resolve conflict between competing representations? This is the question addressed in one way or other in most of the contributions to this volume. On a theoretical or neural level, there are interesting arguments on both sides of the inhibition debate (see Carr & Dagenbach, 1984). However, attempts to address the inhibition question empirically have been marred by difficulties. Very often empirical patterns that look like inhibition, invite an array of non-inhibitory explanations (e.g., MacLeod, this volume). Nevertheless, ultimately it is only through empirical evidence that we will be able to settle the question whether or not cognitive inhibition exists. In the current chapter I will summarize some our own and others' attempts to trace cognitive inhibition in the context of executive selection among competing task sets. I will begin by providing some reasons for why the domain of task-set selection may be particularly well-suited for finding an empirical, behavioral-level signature of inhibition. I will present the lag-2-repetition paradigm (sometimes also referred to as the backward-inhibition paradigm) as the basic experimental tool to trace task-set inhibition as well as some of the basic characteristics of task-set inhibition that have been revealed with help of this paradigm. Then, I will discuss in some length how we can deal with possible alternative, non-inhibitory interpretations. Finally, I will present some new findings form mostly ongoing projects about the neural implementation and associative learning aspects of task-set inhibition.

Why Inhibition during Task-Set Selection?

We can think of a task set as the constellation of attentional settings that allow an individual to perform one among several tasks that are potentially possible in a given situation. For example, a very simple task used in some of our experiments requires subjects to localize a deviant object on a given stimulus dimension (see Figure 1). In this case, the task set would represent the dimension (e.g., color) and ensure an attentional configuration that gives an increased weight to processing information from that dimension. Empirical evidence pointing to the existence of task sets comes from the simple finding that when task sets need to change from one trial to the next, there is a time cost (i.e., the so-called switch cost; Monsell, 2003) irrespective of whether or not specific dimensional values or response parameters change.

Task sets can be triggered in a bottom-up manner. For example, when we step into the kitchen in the morning, we usually do not have to intentionally gear our cognitive system towards brewing coffee and making breakfast. Rather, the general context and environmental cues activate the adequate configuration of processes. While such bottom-up, task-set priming is generally useful, it is also one important reason why we need to be able to exert top-down control on task-set selection. Only through top-down task-set control can we arrive at a stable behavioral pattern in a situation in which the environment triggers multiple possible task sets. For example, because of a looming grant deadline, one may decide to stop all activities such as web surfing and emailing that may otherwise be prompted by sitting in front of the computer. In such situations, a top-down selected set is necessary to maintain a coherent stream of goal-directed activity in the face of constant reminders of other possible activities. Another important reason for top-down controlled task-set control is that sometimes an already firmly established task, and thus fully activated set, needs to be replaced by an alternative task set. Note, that these two different demands on task-set selection, maintenance in the face of competition and the need to flexibly replace a task set by an alternative task set, are difficult to align with each other, a fact that is sometimes referred to as the stability-flexibility dilemma (e.g., Goschke, 2000). To achieve behavioral stability in the face of interference, strongly activated task sets are needed. However, strongly activated task sets should also be particularly difficult to get rid off, once the time for change has come. Theoretically, a mechanism that inhibits the no-longer relevant task may provide a solution to this dilemma. Just as we only dare to speed on our bicycle when there is a functioning break in place, the cognitive system may be able to afford full task-set endorsement only because inhibition allows clearing the slate when a new task set becomes necessary. Occam's razor may advise against invoking extra mechanisms, such as inhibition. However, arguably, if there is a place for a special, inhibitory mechanism, flexible task set control may be a prime candidate to look for it.

A second reason why task-set inhibition may be a particularly promising domain to look for cognitive inhibition is related to the methodological challenge of finding unambiguous evidence for inhibition. Evidence for inhibition typically comes in the form of some kind of change in performance (usually a decrement) in response to a stimulus that at some earlier time was supposedly a target of inhibition (e.g., Tipper & Canston, 1985). However, data patterns of this general type are open to non-inhibitory explanations based on the assumption of proactive interference between ways of dealing with a certain stimulus in the recent past (e.g., not to respond to it) and the currently required way of dealing with that same stimulus (e.g., to respond to it; Neill, 1997). Importantly, a hallmark of such memory influences is their specificity. That is, such influences should be the greater the larger the similarity between the initial suppression/encoding situation and the response/retrieval situations. Standard inhibitory paradigms allow little variation of stimulus aspects, simply because inhibition is usually supposed to be targeted at a specific stimulus representation. However, task sets are per definition abstract and therefore not bound to any specific stimulus situation. This opens up much needed degrees of freedom to manipulate the similarity between past and present selection instances. Later in the chapter, I will show how this freedom allows us to establish inhibitory phenomena that are less susceptible to the proactive-interference argument.

The Lag-2 Repetition Paradigm

How can we find an empirical indication of inhibition on the level of task sets? The basic idea is very simple and rests on two assumptions: First, abandoning a task set in order to switch to a new one goes along with inhibition of the abandoned task set. Second, inhibition renders the task-set representation less available for some time. If this is the case then it should be relatively difficult to return to a recently abandoned task set. Figure 1 shows how we implemented this basic idea using the deviant-detection paradigm introduced earlier. For each trial, one task of three or four possible tasks is cued using a verbal task label. The critical lag-2 repetition sequence is one in which subjects might be cued from the color to the orientation task set and then back to color. Theoretically, inhibition should be needed to go from color to orientation. When going back to color, this task set should be still inhibited and thus more difficult to select compared to a control sequence in which the third-position task is one that was less recently abandoned (see Figure 1). In most implementations of this paradigm, lag-2 repetition and lag-2 change sequences were established by selecting tasks randomly on a trial-by-trial basis and then categorizing task triples post hoc into the critical categories. In the initial paper that introduced this paradigm, results did, in fact, show a slowing of RTs on the third trial of a lag-2 repetition sequence relative to control sequences in the order of 20 to 50 ms (Mayr & Keele, 2000). Subsequently, this basic result has been replicated numerous times across different labs (e.g., Arbuthnott & Frank, 2000; Huebner, Schuch & Koch, 2003, Dreher & Berman, 2002; Mayr, 2002) suggesting that it is a robust phenomenon in task-selection situations. Certainly, an interesting aspect of this phenomenon is that at least without additional assumptions, activation-only theories should have a hard time explaining why activating a representation should render it less accessible for subsequent selection. In contrast, this is exactly what one would expect if inhibition is critical during task-set selection. However, before accepting that this effect is in fact associated with inhibition, we need to explore its characteristics to see whether or not it behaves as we would expect from an inhibitory process that is in the service of flexible task-set control.

Characteristics of Lag-2 Repetition Cost

The initial demonstration of the lag-2 repetition cost in Mayr and Keele (2000) was with simple, perceptual sets. This may raise the question to what degree this is a general effect that can also be found for more complex task sets, or sets that are not bound to perceptual dimensions. In recent years, the effect has been found for more complex perceptual tasks (Mayr, 2001), tasks requiring semantic judgments about digits (e.g., Schuch & Koch, 2003), and tasks that require applying different spatial response rules to invariant stimulus aspects (Mayr, 2002; see also the contribution of Dagenbach and colleagues to this volume). Thus, currently, there is little reason to assume that it is bound to certain stimulus or response characteristics, or certain domains of processing.

Theoretically an even more important question is to what degree the lag-2 repetition cost is actually associated with abstract representations, rather than with specific stimulus characteristics. Mayr and Keele (2000) addressed this question using the deviant-localization paradigm. For each possible dimension, deviations from the standard object, a stationary blue vertical line, could occur in two ways. For example, the object could be tilted either 45 degree to the left or to the right with regard to the orientation dimension, or it could be dyed either cyan or purple on the color dimension. This allowed us to test to what degree the lag-2 repetition effect survived lag-2 changes in the dimensional value. Across five different experiments, we found no statistically reliable influence of value changes on the lag-2 repetition effect. Similarly, we could look at the degree to which lag-2 repetitions versus changes of locations/responses affected the lag-2 repetition effect. For example, on a lag-2 repetition of the color task, the critical color deviant might occur in the upper-left corner on the lag-2 trial and eitehr on the same location or on a different location on the probe trial. Again, we found no statistically reliable influence of this variable. If anything, there was a numerical reduction of the lag-2 repetition cost in case of lag-2 response repetitions. We will come back to the possible role of lag-2 response repetitions when we discuss possible alternative interpretations of the lag-2 repetition cost. For now, the fact that this effect was not reliably modulated by repetitions versus changes of specific stimulus/response aspects suggests that it is in fact bound to the level of abstract task sets rather than to the level of stimulus-specific implementations of a particular task.

A core assumption is that task-set inhibition serves top-down selection of task sets. In other words, we should see the lag-2 repetition cost only when subjects are actually selecting task sets in a top-down manner, but not when tasks are primed bottom-up. Mayr and Keele (2000) tested this assumption, using a variant of the deviant-detection paradigm. Again, subjects had to locate deviants from three possible dimensions (color, orientation, movement). Each display contained one deviant on one of these dimensions. Different from the preceding experiments, the distractor deviant was of a different kind, namely size. In each display, there one rectangle was substantially larger than the rectangle object. However, object size was never one of the cued dimensions (see Figure 2). This setup allowed the comparison of two different conditions. In the top-down condition, subjects received valid precues about the next relevant dimension, just as in the preceding experiments. In the bottom-up condition, they only saw non-informative letter strings as cues. Given that there was only one deviant from one of the three "legal" dimensions in each display, the adequate response could be selected even without prior task-set information. However, our hope was that inclusion of the irrelevant size deviant would make response selection sufficiently hard to motivate subjects to use the precue in the top-down condition. In fact, we found that overall, subjects performed reliably faster in the top-down than in the bottom-up condition, suggesting that they in fact did make use of the task cues. As predicted, there was no lag-2 repetition cost in the bottom-up condition. However, there was a reliable, although very small (7 ms) lag-2 repetition cost. The small size of this effect is not too surprising given that there was really no need for subjects to use top-down control. In fact, we can use for each individual subject the difference score between the top-down and the bottom-up condition is an indicator of top-down control effort. Figure 3 shows the scatterplot underlying the correlation between this variable and the lag-2 repetition cost. As evident, for subjects with large top-down selection costs, the lag-2 repetition cost was substantial, but was basically non-existent for subjects with zero or negative top-down effect. Thus, just as we would expect if task-set inhibition aides between-task conflict resolution during top-down task-set selection, the lag-2 repetition cost only occurred when subjects actually engage in top-down task-set control.

Does the fact that the lag-2 repetition cost is tied to top-down task selection necessarily imply that it is also under top-down control? In their Experiment 4, Mayr and Keele (2000) used, instead of the random cuing procedure, a procedure in which subjects were pre-informed about a sequence of four tasks. Specifically, at the beginning of each new sequence of four tasks, they saw a precue listing the entire sequence of four tasks. After studying this precue, the information disappeared and four displays were presented one at a time and subjects had to execute a response to each one according to the currently relevant, memorized task. The first task in the sequence was a filler task (always the size task), whereas in the sequence positions 2 to 4 the lag-2 repetition versus control sequence was implemented (e.g., color – orientation – color versus movement – orientation – color). Thus, in this case subjects had foreknowledge about the entire sequence, including that when switching in an inhibition sequence from the 2nd to the 3rd task they would inhibit a task that will become relevant again on position 4. Nevertheless, we obtained a very robust lag-2 repetition effect in this situation. In other words, it seems that task-set inhibition is not affected by foreknowledge and thus is not under top-down control. More likely, task-set inhibition is a mandatory process that may be recruited when a top-down initiated switch between tasks occurs.

If task-set inhibition aids the resolution of between-task conflict, it should be particularly strong when between-task competition is high. Gade and Koch (in press, see also Mayr & Keele, 2000, Exp. 2) provided interesting evidence in this regard by manipulating the temporal relationships between successive trials. Note, that for a specific inhibition (or control) sequence, there are two potential transitions: the one between the lag-2 and the lag-1 of the sequence and the one between the lag-1 and the probe trial of the sequence. Inhibition should arise at the transition between the lag-2 and the lag-1 trial. Arguably, inhibition should be particularly important when the task for the lag-1 trial needs to be initiated when the lag-2 task is still highly active. The authors manipulated this aspect by varying the time between the response to the previous task and the cue to the next task randomly on a trial to trial basis (while keeping the total time between trials constant; see Meiran, 1996). As expected, inhibition was much stronger for short response-to-cue intervals (i.e., when the preceding task should still be highly active) than for long intervals. Importantly, inhibition was only dependent on the response-to-cue interval between the lag-2 and the lag-1 trial of the sequence (i.e., where inhibition should occur) but not to the interval between the lag-1 and the probe trial of the sequence. This result is highly consistent with the notion that inhibition is recruited in response to between-task inhibition.

Another recent result seems to provide additional, more specific information about the type of competition that elicits task-set inhibition. Schuch and Koch (2003) used the basic lag-2 repetition situation, but combined it with a nogo paradigm. On 25% of the trials, a tone was sounded along with the stimulus, indicating that subjects should not execute the response. The critical comparison was between inhibition sequences for which the lag-1 trial task happened to be a nogo trial and sequences for which all trials were go trials. In case of a lag-1 nogo trial, subjects had to prepare for the task indicated by the cue for the lag-1 trial, but then had to refrain from executing the relevant response selection process. If inhibition arises from cue-driven preparation alone, then we should see inhibition irrespective of the status of the lag-1 trial. However, if inhibition is in some way linked to the actual response-selection process then we should see inhibition only for go trial sequences. In fact, robust lag-2 repetition costs were obtained after go trials, but zero costs were found after nogo trials. These results suggest that cue-driven processes may be necessary, but not sufficient to elicit inhibition. Rather, only once a task set is applied to a particular stimulus situation, is inhibition triggered. While an interesting and surprisingly clean result, there are some open questions. For example, it is not easy to reconcile the fact that inhibition was highly sensitive to the response-cue interval in the Koch and Gade (in press) study when according to Koch and Schuch (2003) the critical competition takes place only after stimulus presentation. One possibility is that with 25% nogo trials, subjects may encode the cue, but not actually use the information to prepare the next task until after they could be certain that a response would be required (e.g., Kleinsorge & Gajewski, 2004). Note, when Mayr and Keele (2000) contrasted the role of task cues they found that inhibition was limited to subjects who actually made use of the cue (Figures 2 and 3). In so far unpublished work, we have constructed conditions in which there was a greater incentive for preparation even in case of nogo trials. At least in some of these situations we have found reliable lag-2 repetition costs. Overall, we believe that Schuch and Koch's results suggest that response selection may be one important triggering condition for task-set inhibition. However, the important question to what degree it is the only possible trigger, or whether mere thinking of an alternative task can also elicit inhibition, requires additional scrutiny.

While it is currently not 100% clear what exactly triggers task-set inhibition, there is some information available about the level of representation that is the most likely targeted by inhibition. Several authors have argued for a kind of 2-component view of task selection (e.g., Mayr & Kliegl, 2003; Rubinstein, Evans, & Meyer, 2001). The first component involves retrieving or activating a symbol-level goal in working memory (e.g., "do the color task"). This may be the most important aspect of what can be achieved after the cue has been presented, but before appearance of the stimulus (i.e., the preparation interval). The second component is the actual configuration that results from applying this goal to a particular stimulus situation. It is currently not clear how much of this component can be actually prepared. However, it seems that it is this component that is targeted through task-set inhibition. One critical result in this regard comes from a study by Mayr and Kliegl (2003). They used a paradigm in which two different cues where mapped onto each of three tasks. This allowed the implementation of lag-2 task repetitions in which either the cues stayed the same or changed. In other experiments in this paper, the authors had demonstrated that a large part of switch costs between tasks arose from a mere change in cue and that this component is specifically sensitive to the cue-stimulus interval. From this, the authors had concluded that the cue-associated switch cost reflects the first, the symbol level component, whereas the smaller switch-cost component associated with the actual switch in task reflects the change in configuration. Using this paradigm with three different tasks, the authors looked at lag-2 repetition transitions with or without cue changes. Interestingly, a robust lag-2 cost was found for cue changes. This suggests that inhibition is associated with the actual representation that is used to operate on a particular stimulus in task-congruent manner.

A similar conclusion can be derived from results regarding the dependency of the lag-2 repetition costs on the cue-stimulus interval. If the representation that undergoes inhibition is easily targeted through preparatory activity, then the lag-2 repetition cost can be expected to decrease when there is sufficient time to prepare. However, studies looking at this aspect have regularly found the lag-2 repetition cost to be completely insensitive to this variable (e.g., Mayr & Keele, 2000; Shuch & Koch, 2003). Again, this suggests that inhibition targets the actual configuration that is applied to the stimulus. It seems possible that this representation, because it is so closely tied to the stimulus becomes activated only when the stimulus appears and therefore is out of reach of purely cue-driven preparatory processes (for additional discussion of this aspect see Mayr & Keele, 2000).

The results we have reviewed here suggest that lag-2 repetition costs behave as we would expect from a process that is tied to top-down task selection and that is geared towards resolving between-task competition. It occurs only when top-down control is required and it is sensitive to the amount of between-task competition (as indexed through the temporal interval between successive tasks). While being tied to top-down control, it is not itself penetrable through top-down control. It occurs even when this implies inhibiting a task that will be relevant again in the very near future and once a task is inhibited, preparation does not allow reversing the inhibition associated with that task. At the same time, there remain important open questions regarding the sufficient and necessary conditions for inhibition. Specifically, it will be important to firmly establish whether it is it enough to think about a new task, or whether one has to execute a new task in order to accomplish inhibition of the preceding task.

Non-inhibitory Explanations of the Lag-2 Repetition Cost

There are three different types of non-inhibitory explanations of the lag-2 repetition cost that need to be considered: 1) sequential expectancies, 2) activation of likely successor tasks, and 3) instance-based memory effects. Whereas the first two are specific to the lag-2 repetition effect, the third explanation is of more general importance, because it has been also applied rather successfully in the context of other inhibitory paradigms.

Sequential Expectancies

When people are asked to make random, sequential choices, they usually deviate in a predictable manner from randomness. For example, in case of binary choices they tend to produce a higher rate of alternations than random choice would imply (i.e., the gambler's fallacy). The typical explanation for this behavior is that people use a representativeness-within-small-runs heuristic. In other words, they try to produce short sequences that "look" representative. In binary-choice cases, this would favor a sequence such as ABA over AAA, even though both triples are of course equally likely. Applying this to the lag-2 repetition paradigm where subjects need to select among three different tasks one might expect them to favor an ABC over an ABA sequence, simply because only the first kind involves all possible tasks and therefore may seem more representative. If subjects actually entertain such sequential expectancies, the lag-2 repetition effect may arise from an expectancy violation rather than task-set inhibition. However, there are two results that speak against such an interpretation of the leg-2 repetition effect, both of them already mentioned in the preceding section. The first is that the inhibition effect is not affected by the duration of the preparatory interval. If the lag-2 repetition effect was driven by a sequential expectancy, why would such an expectancy not be overruled, or at least attenuated, by a cue-induced 100% valid expectancy about the task to come. The same argument holds for the second relevant result, namely that the lag-2 repetition cost is found even when subjects have complete foreknowledge about an upcoming sequence of tasks. It seems hard to explain how sequential expectancies could develop about a sequence of tasks for which a complete plan exists.

Activation of Likely Successor Tasks

Usually, the lag-2 repetition cost is assessed in a situation with 100% task switches. This is done to maximize the occurrence of the relevant inhibition or control triples. However, it is possible that this may lead subjects to actively prepare after each task (e.g., color) for the two, now potentially relevant tasks (orientation or movement). Thus, both tasks are now preactivated, but only one of them, let us assume the movement task, will be actually used on next trial. The orientation task, having been prepared but still unused, may continue to be residually active and this could give this task a slight edge in case it needs to be selected on the following trial (which would serve as probe trial, of a lag-2 non-repetition sequence). As a result, performance on the third trial of a control sequence (color-movement-orientation) might be faster than the third trial of a lag-2 repetition sequence (color-movement-color). Obviously, this explanation of the lag-2 repetition cost requires no assumption about inhibition. However, just as the sequential expectancy account it is also based on expectancies. Therefore all the counterarguments leveled against the sequential expectancies account are also relevant here. In addition, there are two other results that speak against this particular account. First, if this account were correct, one would expect that the lag-2 repetition cost should be particularly large when there is plenty of time between the response to the lag-2 trial and the task cue for the lag-1 trial so that the critical activation pattern can develop. However, as we had reported in the previous section, the lag-2 repetition cost is at its maximum precisely when the time between the response to the lag-2 trial and the cue to the lag-1 trial is very short (Gade & Koch, in press). Second, there are a number of studies in which transition probabilities between tasks were balanced (e.g., Arbuthnott & Frank, 2000; Dreher & Berman, 2002; Mayr & Keele, 2000). In such a situation, there is little reason to assume that subjects prepare for task changes. Nevertheless, usually robust lag-2 repetition costs can be found even in these situations.

Instance-Based Episodic Priming

This is in many ways the most interesting and probably also most challenging non-inhibitory account. It is interesting, both because it arises from a well-founded theory on the interplay between attention and memory (i.e., instance theory; Logan, 1988) and because versions of it have been applied to other inhibitory phenomena (e.g., Neill, this volume). The account basically holds that each selection episode is encoded into memory as a *specific* instance. Such "memory snapshots" may contain all selection-relevant parameters, such as the stimulus location, the values on the relevant dimension and possibly also the locations and values of distractor objects. When, on one of the following trials, the stimulus situation happens to match in some way the encoded episode then it will be automatically retrieved. In some cases, automatic retrieval might help performance on this trial, in particular if there is a complete match between the encoded selection parameters and the currently required selection parameters. At other times, when the match is incomplete, a mismatch cost may arise. For example, as mentioned in the Introduction, the negative-priming effect can be explained in terms of episodic, instance-based priming: When subjects try to ignore a distractor object on the prime trial, the resulting memory trace may include a nogo-tag that becomes associated with the representation of that object. This memory trace would then produce proactive interference when, on the probe trial, a response has to be made to that same object. Results with the negative-priming paradigm have demonstrated that episodic priming effects are clearly one important factor (see Neill, this volume).

As an aside, it is interesting to note that from an inhibition-theory standpoint, inhibition is exactly what might help against unwanted intrusions from earlier selection episodes. This may have interesting consequences when comparing people who presumably differ in inhibitory control ability (e.g., old and young adults). People with good inhibitory control may be able to remain relatively unaffected by episodic priming effects and therefore negative priming (or some other inhibitory phenomenon) may actually represent "true" inhibition. In contrast, people with low inhibitory control may be particularly vulnerable to proactive interference from earlier selection instances. They may therefore show the inhibitory phenomenon, but for the wrong reasons--ironically exactly because their inhibitory control is low.

The episodic priming account is a viable competitor for the inhibitory explanation of the lag-2 repetition cost. Let us assume the lag-2 trial of an inhibition sequence requires locating the color deviant (e.g., a pink rectangle in the upper right corner).

According to the episodic priming account, all of these aspects are encoded in the memory instance. If on the probe trial, the color task is cued again, the cue may also lead to reactivation of this particular memory instance. However, now the critical deviant may be a purple rectangle in the lower left corner. Thus, there is a mismatch between the required selection setting and the earlier encoded memory instance and it is this mismatch that might cause the lag-2 cost.

Whereas the episodic-priming account rendered the negative-priming paradigm basically useless as a tool to study inhibitory control, there are reasons to believe that the lag-2 priming paradigm may be more resistant to this particular alternative explanation. Given that in this paradigm, inhibition targets the task-set level rather than the level of specific stimulus or response representations, we are in an excellent position to vary the similarity of stimulus-response aspects across the critical lag-2 repetitions. This is important, because specificity is an important factor that should modulate episodic priming effects. Thus, if the lag-2 repetition cost actually arises because of the mismatch between a memory instance and the current selection demands then we should see the cost disappear, and even turn into a benefit, if we establish conditions in which the memory instance and the selection needs match.

Mayr and Keele (2000) tried to do exactly that by comparing cases in which potentially critical parameters matched between lag-2 trial and the probe trial of a lag-2 repetition sequence with cases where there was a mismatch. With regard to the feature value and the location/response aspect there were no reliable effects. This lack of a clear mismatch effect is not entirely reassuring because these analyses could focus only on one potentially relevant parameter at a time. It is possible that only a complete match on all parameters would produce a priming benefit and that all other cases might produce a cost (e.g., Hommel, 2004). The deviant localization paradigm simply involved too much stimulus variability to allow an analysis of all potentially relevant parameters, including their interactions as well as the most informative case of a complete match on all aspects. Mayr (2002) therefore used a paradigm in which stimulus variability was greatly reduced. Subjects saw on each trial a circle appearing in one of four corners of a square. Responses were determined by three different spatial transformation rules (horizontal, vertical, and diagonal). For example, for the horizontal rule, a stimulus in the upper left corner would require a horizontal translation and thus a response in the upper right corner. Given that there were only four different stimulus positions, 25% of all lag-2 repetition cases involved a complete stimulus match. If episodic priming is in fact responsible for the lag-2 repetition cost, then we should see a benefit for these trials. However, fortunately for the inhibition view, there was a lag-2 repetition cost even in the case of these complete matches. Of course this result does not speak against the reality of instance-based priming effects in general (see Neill, this volume). However, together with the results reported in Mayr and Keele (2000; see also Schuch & Koch, 2003), there is this far no evidence that they are behind the lag-2 repetition cost.

To summarize, the lag-2 repetition cost cannot be accounted for in terms of sequential expectancies, activation of likely competitors, or instance-based priming. This strengthens the conclusion that it the lag-2 repetition cost does in fact reflect inhibition targeted at no-longer relevant task sets.

Neural-Level Explorations

Aron, Robbins and Poldrack (2004) have recently proposed that an area in the right-inferior frontal cortex is critically involved when inhibitory control is used to resolve conflict. This conclusion is based on two results. First, stop-signal time, that is the time an individual needs to stop an initiated response process, increases linearly and quite orderly with the size of the lesion in this region. With the same group of patients, Aron and colleagues had also looked at task-switching performance (Aron, Monsell, Sahakian, and Robbins; 2004). Here, they found that right-frontal, but not left-frontal patients showed much increased error switch costs and somewhat increased RT switch costs when the switch-trial stimulus was bivalent. They argued that deficient inhibition of the previous task set allowed for an increase in task-set interference after a switch. While these results are certainly consistent with the claim that right-inferior frontal patients have an inhibitory deficit, they are not fully conclusive. The degree to which inhibition proper is involved in the stop-signal procedure is currently not resolved. In fact, the original model simply assumes a race between two independent processes (a go process and a stop process; see Logan, this volume). Similarly, the task-switching result could be easily explained in terms of slowed activation of the currently relevant task, which would give the previously relevant task more time to interfere (see Cohen & Dehaene, 1998). In brain-imaging studies, right prefrontal activations are often found in situations of high response conflict (Garavan, Ross, & Stein, 1999; Hazeltine, Poldrack, & Gabrieli, 2000) and maybe this is simply an area that is responsible for response conflict resolution without any necessary implication for inhibitory control.

In this situation, it would be useful to obtain results with a paradigm that provides a more unambiguous record of inhibitory activity. We therefore compared seven left prefrontal and four right prefrontal patients in a task-switching paradigm that allowed us to assess local switch costs, the global costs that emerge when comparing task-switching performance with single-task performance, and the lag-2 repetition cost. We found that the left-frontal patient showed increased local and global costs. Interestingly, they had increased costs even when stimuli were univalent, suggesting a rather basic deficit in terms of activating the next task set. However, left-frontal patients showed lag-2 repetition costs that were if anything numerically larger than those of control subjects. In contrast, right-frontal patients had normal local and global costs, but no lag-2 inhibition cost. This result is fully consistent with the hypothesis put forward by Aron et al. (2004) that inhibition is associated with right prefrontal cortex. However, there was one aspect that was at odds with the Aron et al. (2004) task-switching study: Whereas their right-frontal patients had a switching deficit, our patients seemed quite unaffected by the switching demands, despite the fact that they did not show inhibition. If inhibition is involved in switching, shouldn't lack of inhibition lead to greater switch costs? A follow-up experiment provided a hint to what may have been going on. Aron at al. had used a paradigm in which subjects worked on the same task for three trials in a row, a situation that allowed strong endorsement of each task. In contrast, to get at the lag-2 repetition effect, we had to use a paradigm in which tasks could change randomly after each trial. Such a procedure probably leads to less endorsement of each individual task and therefore also less need for inhibition to get rid of it on the next trial. In our follow-up experiment, we used a procedure with runs of three tasks between switches (thus promoting full task-set endorsement, but no assessment of lag-2 repetition costs). In this situation, right-frontal patients actually did show a trend towards an increased switch cost that, at least in terms of numerical RT effect, was similar in size to the one found by Aron et al. (2004). This pattern of results suggests that the typical lag-2 repetition paradigm in which tasks change every trial may be well-suited to track task-set inhibition but, because it does not encourage full endorsement of each task, it may not be the best paradigm to reveal the effects of an inhibitory deficit on switch costs.

Together, our results and the findings by Aron and colleagues point to the involvement of right prefrontal cortex in the manifestation of inhibitory control. However, it is also important to point out that so far these data do not tell us what exactly this area is doing. Probably the most straightforward hypothesis is that it directly involved in initiating inhibition. However, a slightly more complicated model might follow from the premise that prefrontal neurons are primarily involved in representing potentially relevant aspects of the current task environment (e.g., Duncan, 2005). Recent neuroimaging findings are of interest in this regard. Using the lag-2 repetition paradigm, Dreher and Berman (2003) obtained a robust increase in right prefrontal activation when subjects had to switch back to the most recently suppressed task set. In other words, right prefrontal cortex seemed involved in expressing the lag-2 repetition cost when returning to a recently abandoned task, but not necessarily at the time when one would expect inhibition to operate (i.e., when switching to Task B in an ABA sequence). One way to interpret these data is by assuming that whereas left-prefrontal cortex implements selection of a single, coherent task set (e.g., Mayr & Kliegl, 2000, 2003), right prefrontal cortex maintains a broader context (e.g., Keele, Ivry, Mayr, Hazeltine, & Heuer, 2003). Such a broader context would allow backtracking or monitoring of success across different tasks in multiple-task performance and it may include a trace of the suppression associated with task sets that had recently been deactivated. Thus, while right prefrontal cortex may represent the fact that a task has been suppressed, it is not necessarily the origin of the suppression. This interpretation is consistent with recent neuroimaging results reported by Braver, Reynolds, and Donaldson (2004) suggesting that during task-switching blocks there is sustained activity in the right-prefrontal cortex, possibly reflecting the increased load of tracking multiple tasks. In contrast, activity associated with local switch transitions is expressed in left-frontal cortex. One novel prediction derived from this view is that deficits associated with right frontal lesions should become particularly apparent in multitask situations in which task selection occurs in a more self-guided manner (i.e., based on a record of previous task selections; Arrington & Logan, 2004) than when selection is under the control of external cues. Whether prefrontal cortex actually houses certain control operations or whether it fulfills strictly representational functions is a fundamental question. At least for the case of inhibition, further research on how exactly right prefrontal cortex contributes to the lag-2 repetition cost should be informative with regard to this important issue.

Limitations and Open Issues

So far, we have emphasized the virtues of the leg-2 repetition cost as an indicator of inhibitory control as well as the things we have learned about inhibition using this indicator. However, as with every paradigm, there are certain shortcomings that are worth pointing out and there are many things that are yet to be learned. In this final section, we highlight some of these issues.

Assessing Aftereffects of Inhibition versus Assessing Inhibition when it Actually Happens

Basically all known behavioral, inhibitory paradigms, including the lag-2 repetition paradigm, try to catch the aftereffect of inhibiting a representation. This makes it difficult to study inhibition as it actually occurs. For example, it is exactly for this reason that the exact point during a task switch at which inhibition comes into play has remained elusive (i.e., during the processing of the cue for the lag-1 trial task or during processing the lag-1 trial stimulus). Similarly, the results of Dreher and Berman (2002), regarding right prefrontal activity associated with the lag-2 repetition cost leave open the question of which brain area is involved with the inhibitory process itself.

Unfortunately, similar ambiguities arise when trying to infer from the lack of a group difference in lag-2 repetition cost that there is no group difference in inhibition. For example, Mayr (2001) found no evidence for a reduced lag-2 repetition cost in old adults, even though a prominent theory states that inhibition is generally reduced in old age (see Hasher, this volume). The Mayr (2001) result could be interpreted as being inconsistent with the inhibition-deficit view (Burke, this volume). However, given that we cannot observe the inhibitory process itself, but only the aftereffect of this process, it is also possible that while the endproduct may be equivalent for old and young adults, it may be much harder (i.e., take much more time) for older adults to arrive at this endproduct. Thus, future progress in this area hinges on coming up with a paradigm that allows a better way of tracking what is going on during inhibition of a currently unwanted task set. One interesting path that we are currently exploring is the use of an analog of Anderson's think/no-thing paradigm (Anderson & Green, 2001; Anderson, this volume) in the task selection domain. This paradigm has proven useful in establishing neural correlates of memory inhibition as it occurs and relating these to the behavioral aftereffects of inhibition (Anderson, Ochsner, Kuhl, Cooper, Robertson et al., 2004).

Inhibition and Learning?

We have shown that task-set inhibition itself, at least as indexed by the lag-2 repetition cost, is not subject to top-down control. However, this does not mean that inhibition is immune to contextual factors. All studies using the lag-2 repetition paradigm reviewed this far, used either a 100% switch rate or equiprobable transitions between all possible tasks (leading to a 66.6% switch rate for three tasks). Philip and Koch (in press) recently used lower switch rates (50%) and found that the lag-2 repetition cost disappears. One interpretation of this result is that the amount of inhibition exerted during task switches is modulated through the general switch rate. In this respect, the amount of inhibition may be a general control parameter that is adjusted according to current context (Philip & Koch, 2005). However, there is another interesting possibility: Maybe it is not so much the general control context that is critical here, but rather specific transition probabilities between pairs of tasks.

Let us consider a situation with a 100% switch rate. In such a situation, subject might learn that a particular task (e.g., A) can be followed by all other possible tasks (A->B, A->C), but never by itself. Thus, possibly individual, task-specific inhibitory links are established (i.e., from task A to task A) which are ultimately responsible for the lag-2 repetition cost. This model makes two unique predictions: First, if we manipulate between-task transition probabilities such that only particular tasks are never repeated then we should see lag-2 repetition costs only for these tasks. Second, the ability to establish inhibitory links between task sets should not only function in a self-inhibitory manner. Rather, inhibitory association should evolve between any pair of tasks that never (or only rarely) follow each other. For example, we might construct transition probabilities such that a specific task (e.g., A) is never followed by a specific other task (e.g., B). During the transition from A to a "legal" successor task (e.g., C), task B is inhibited. If then on the next trial, B happens to become relevant, then it should still be in an inhibited state and thus difficult to active. In recent, and in part still ongoing work, we have found evidence supporting each of these novel predictions.

At first sight, these results seem to suggest a radical alternative to the idea we had put forward in our Introduction, namely that task-set inhibition is an executive control operation that is generally applied when switching from one task to the next. Rather, task-set inhibition may be the result of specific learning experiences that help to sculpture the internal representation of the task space according to current task demands. However, given that this work is relatively recent, it is too early to draw firm conclusions. For example, the fact that we find that specific transition probabilities modulate the development of inhibitory effects does not rule out the possibility that general control settings (as proposed by Philip and Koch) may also play a role. In addition, we had already reviewed results suggesting that lag-2 repetition costs can be observed when the transition to each possible task (including task repetitions) is equally likely (e.g., Mayr 2002). If inhibition were exclusively tied to unlikely transitions, there would be no reason to expect lag-2 repetition costs in case of equiprobable transitions. Thus, so far, the existing evidence seems most compatible with the view that lag-2 repetition costs may reflect a general inhibitory process that is modulated by specific (and possibly also general) control demands. Clearly though, in future work it will be important to further examine the role of specific and general experiences in shaping inhibitory control.

Conclusion

In this chapter I have focused on reviewing the range of empirical findings resulting from the lag-2 repetition paradigm. Despite its fairly recent entry as an experimental paradigm, it has already provided a number of important empirical results about the nature of task-set inhibition. As an answer to the main underlying question of this volume – is there cognitive inhibition? – these results suggest a cautious "yes". In many ways, the lag-2 repetition cost behaves as one would expect from a mechanism involved in resolving between-task competition and so far, it has proven relatively robust in the face of potential alternative interpretations. At the same time, both the neural-level findings and, in particular, the novel results regarding the relationship between associative-learning and inhibition suggest new avenues of research that promise important insights about the nature of cognitive inhibition.

References

Allport, D. A., Styles, E. A., & Hsieh, S. (1994). Shifting intentional set: Exploring the dynamic control of tasks. In C. Umilta & M. Moscovitch (Eds.), Attention and Performance XV, (pp. 421-452). Hillsdale, NJ: Erlbaum.

Anderson, M. C., & Green, C. (2001). Suppressing unwanted memories by executive control. *Nature, 410*, 131–134.

Anderson, M.C., Ochsner, K.N., Kuhl, B., Cooper, J., Robertson, E., et al. (2004). Neural systems underlying the suppression of unwanted memories. *Science, 303*, 232-235.

Arbuthnott, K. D. & Frank, J. (2000). Trail Making Test, Part B as a measure of executive control: Validation using a set-switching paradigm. *Journal of Clincal and Experimental Neuropsychology, 4*, 518-528.

Gade, M. & Koch, I. (in press). Linking inhibition to activation in the control of task sequences. *Psychonomic Bulletin & Review*.

Aron, A. R., Robbins, T. W., & Poldrack, R. A. (2004). Inhibition and the right inferior frontal cortex. *Trends in Cognitive Sciences*, *8*, 2004.

Aron, A. R., Monsell, S., Sahakian, B. J., & Robbins, T. W. (2004). A componential analysis of task-switching deficits associated with lesions of left and right frontal cortex. *Brain*, 127, 1561-1573.

Arrington, C. M., & Logan, G. D. (2004). The cost of a voluntary task switch. *Psychological Science, 15*, 610-615.

Braver, T., Reynolds, J.R., & Donaldson, D. I. (2003). Neural mechanisms of transient and sustained cognitive control during task switching. *Neuron, 39*, 713-726.

Cohen, L., & Dehaene, S. (1998). Competition between past and present: Assessment and interpretation of verbal perseverations. *Brain, 121*, 1641-1659.

Dagenbach, D. & Carr T.H. (1984). *Inhibitory processes in attention, memory, and language.* San Diego, CA: Academic Press.

Dehaene, S., Kerszberg M., & Changeux, J.-P. (1999). A neuronal model of a global workspace in effortful cognitive tasks. *Proceedings of the national Academy of Sciences, 95*, 14529-14534.

Dreher, J.-C., & Berman, K. F. (2002). Fractionating the neural substrate of cognitive control processes. *Proceedings of the National Academy of Sciences*, *99*, 14595-14600.

Garavan, H., Ross, T.J., & Stein E. A. (1999). Right hemispheric dominance of inhibitory control: An event-related functional MRI study. *Proceedings of the National Academy of Sciences, 96*, 8301-06.

Goschke, T. (2000). Intentional configuration and involuntary persistence in task set switching. In S. Monsell & J. Driver (Eds.), *Control of cognitive processes: Attention and Performance XVIII.* Cambridge, MA: MIT Press.

Hazeltine, E., Poldrack, R., & Gabrieli, J.D.E (2000). Neural activation during response competition. *Journal of Cognitive Neuroscience, 12*, 2000, 118-129.

Hübner, M., Dreisbach, G., Haider, H., & Kluwse, R.H. (2003). Backward inhibition as a means of sequential task control: Evidence for reduction of task competition. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 29*, 289-297.

Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgement of representativeness. *Cognitive Psychology, 3,* 430-454.

Keele, S., Ivry, R., Mayr, U., Hazeltine, E., & Heuer, H. (2003). The cognitive and neural architecture of sequence representation. *Psychological Review*, *110*, 316-339.

Kimberg, D. Y., & Farah, M. J. (1993). A unified account of cognitive impairments following frontal lobe damage: The role of working memory in complex, organized behavior. *Journal of Experimental Psychology: General, 122,* 411-428.

Kleinsorge, T., & Gajewski, P.D. (2004). Preparation for a forthcoming task is sufficient to produce subsequent shift costs. *Psychonomic Bulletin and Review, 11,* 302-306.

Koch, I., Gade, M., & Philipp, A. M. (2004). Inhibition of response mode in task switching. *Experimental Psychology, 51*, 52-58.

Mayr, U. (2001). Age differences in the selection of mental sets: The role of inhibition, stimulus ambiguity, and response-set overlap. *Psychology and Aging, 16,* 96-109.

Mayr, U. (2002). Inhibition of action rules. *Psychonomic Bulletin and Review*, *9*, 93-99.

Mayr, U., & Keele, S. W. (2000). Changing internal constraints on action: The role of backward inhibition. *Journal of Experimental Psychology: General, 129*, 4-26.

Mayr. U. & Kliegl, R. (2003). Differential effects of cue changes and task changes on task-set selection costs. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29,* 362-372.

Mayr, U., Diedrichsen, J., Ivry, R., & Keele, S., (in press). Dissociating task-set selection from task-set inhibition in prefrontal cortex. *Journal of Cognitive Neuroscience*.

Meiran, N. (1996). Reconfiguration of processing mode prior to task performance. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22,* 1423-1442.

Neill, W. T. (1997). Episodic retrieval in negative priming and repetition priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23,* 1291-1305.

Philipp, A. M. & Koch, I. (in press). The relation of task inhibition and task repetition in task switching. *European Journal of Cognitive Psychology.*

Rogers, R. D., & Monsell, S. (1995). The cost of a predictable switch between simple cognitive tasks. *Journal of Experimental Psychology: General, 124*, 207-231.

Rubinstein, J., Meyer, E. D., & Evans, J. E. (2001). Executive control of cognitive processes in task switching. *Journal of Experimental Psychology: Human Perception and Performance, 27,* 763-797.

Schuch, S. & Koch, I. (2003). The role of response selection for inhibition of task sets in task shifting. *Journal of Experimental Psychology: Human Perception and Performance, 29,* 92-105.

Tipper, S. P., & Cranston, M. (1985). Selective attention and priming: Inhibitory and facilitatory effects of ignored primes. *Quarterly Journal of Experimental Psychology, 37,* 591-611.

Figure Captions

Figure 1. Sample displays for a lag-2 repetition and a lag-2 change sequence.

Figure 2. Sample displays for top-down and bottom-up lag-2 repetition sequence using in Mayr and Keele (2000, Experiment 3).

Figure 3. Scatterplot showing individual subjects' lag-2 repetition costs as a function of their topdown control score (i.e., mean RT in top-down condition minus mean RT in bottom-up condition; Mayr and Keele, 2000, Experiment 3).

Figure 1

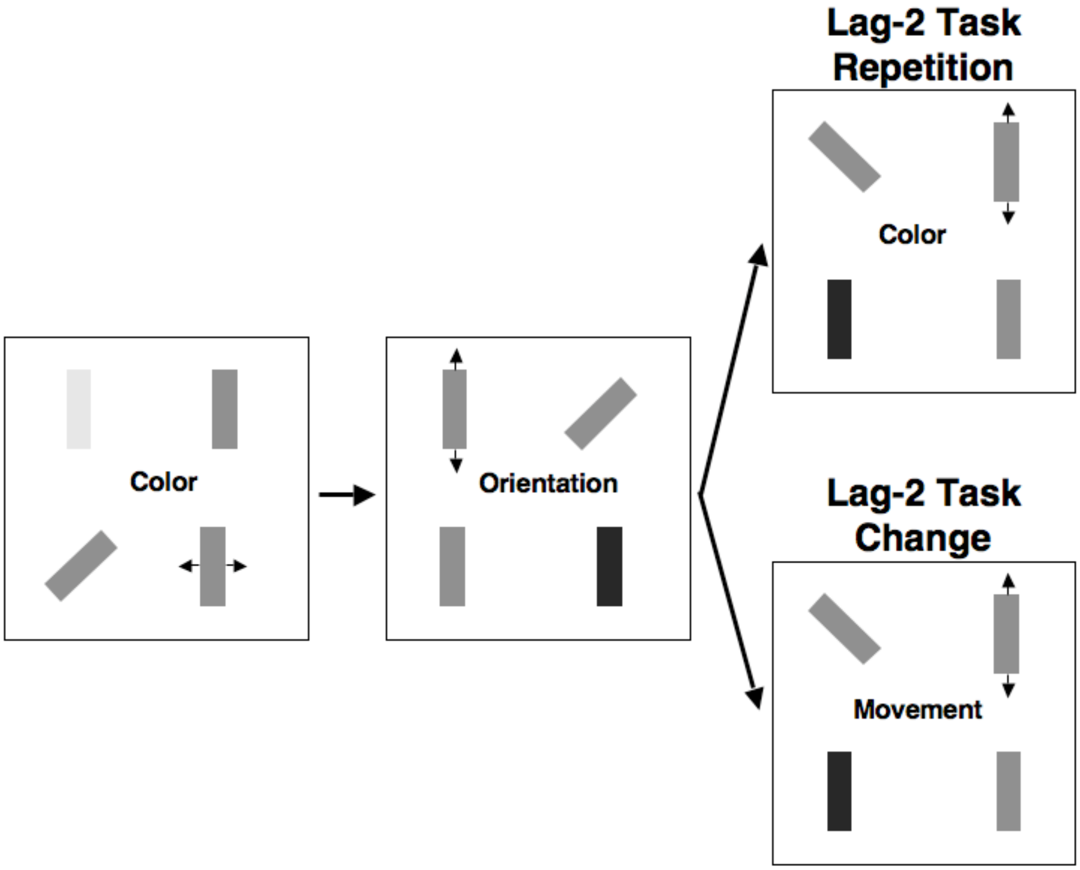
Lag-2 Task Repetition
Lag-2 Task Change
Color
Orientation
Color
Movement

Figure 2

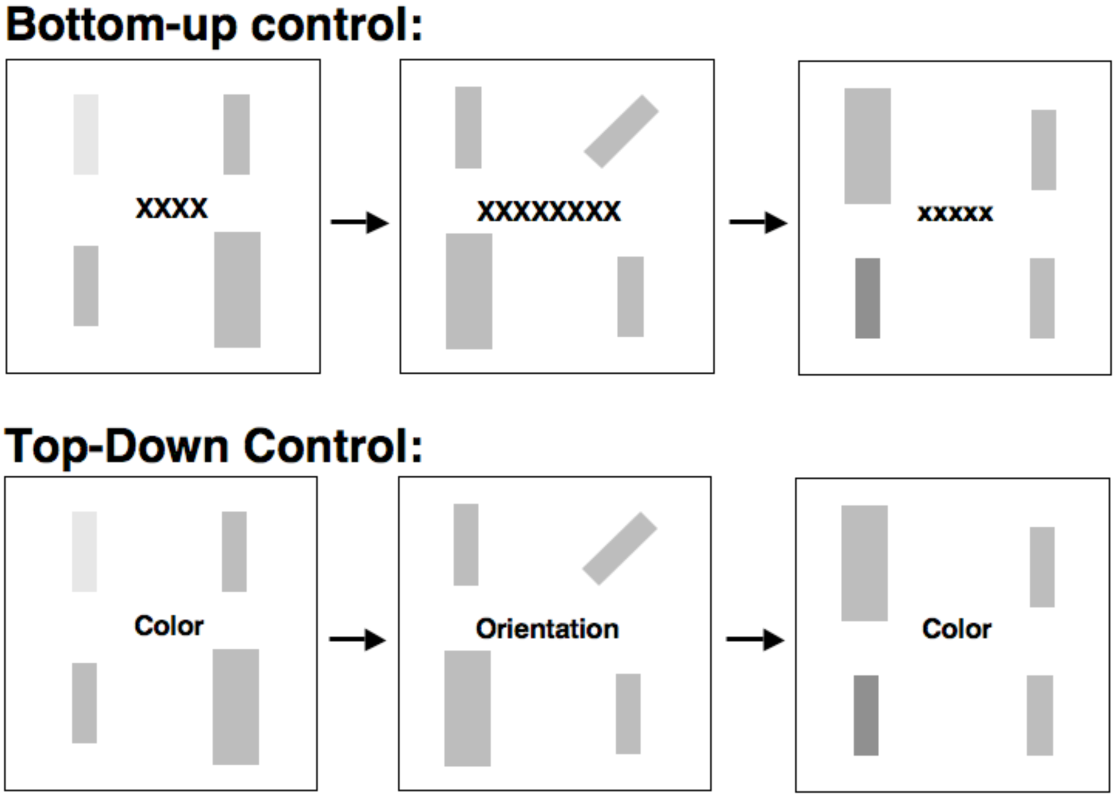
Bottom-up control:
XXXX
XXXXXXXX
xxxxx
Top-Down Control:
Color
Orientation
Color

Figure 3

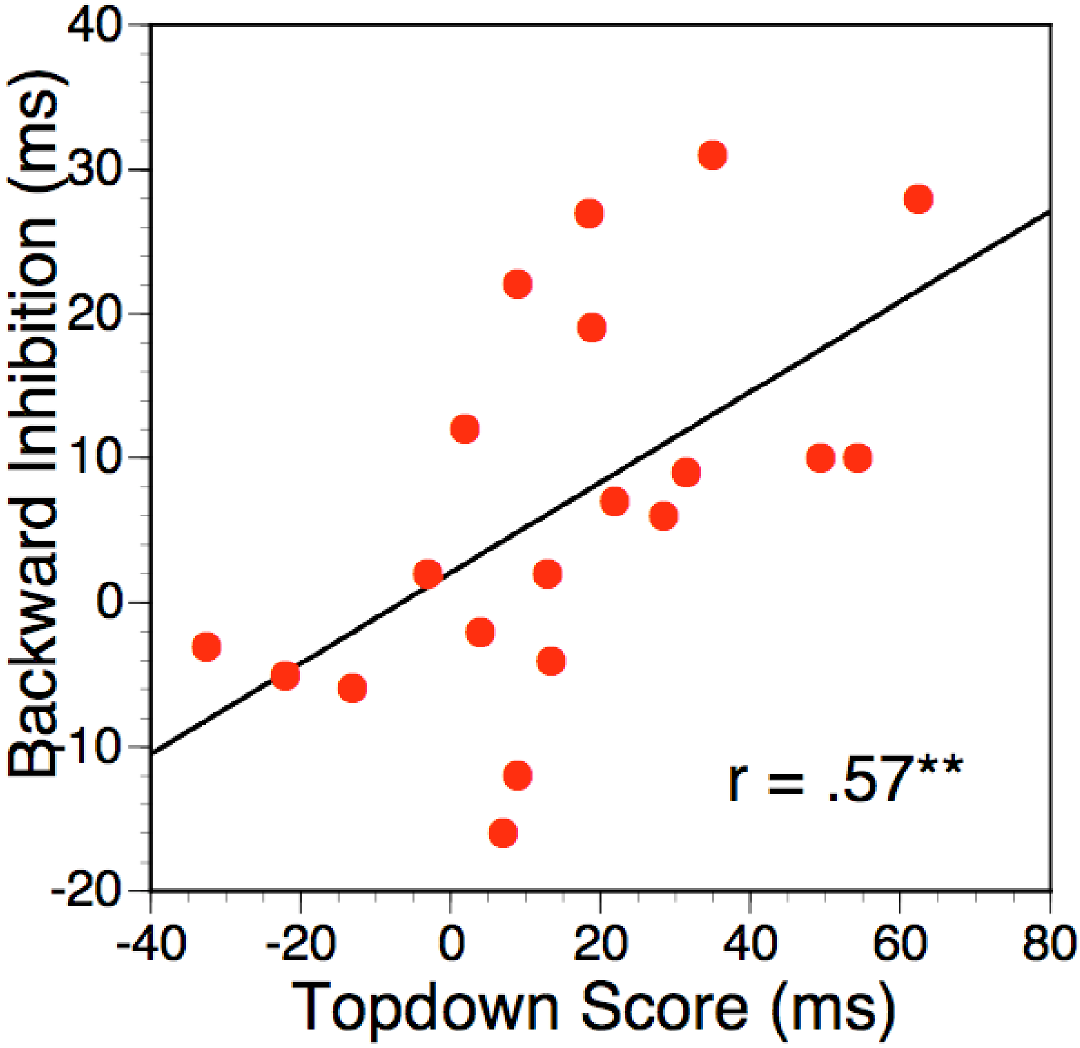

Backward Inhibition (ms)
40
30
20
10
0
-10
-20
-40
-20
0
20
40
60
80
Topdown Score (ms)
r = .57**