

Seminar: Cognitive and Behavioral Neuroscience Seminar (603)
Meeting Date: February 21, 2002

Chairs: Herbert S. Terrace, Peter Balsam, Jon Horvitz, and Yaakov Stern

Speaker: Jonathan D. Cohen

Topic: The Neural Bases of Conflict Monitoring and Cognitive Control: Computational Modeling and Neuroimaging Studies

Seminar Participants:

Herbert S. Terrace, Columbia University, Psychology
Jim Moeller, Columbia University, Psychiatry and Neurology
John Fossella, Cornell Medical School
Matt Davidson, Sackler Institute, Cornell Medical School
Yael Cycowicz, New York State Psychiatric Institute
Helen Gaeta, New York State Psychiatric Institute
David Friedman, New York State Psychiatric Institute
Colin Beer, Rutgers University
SH Lisanby, New York State Psychiatric Institute, Psychiatry
Matthias Tabert, New York State Psychiatric Institute, Psychiatry
Leann Kinnuren, New York State Psychiatric Institute, Psychiatry
Hans Lou, New York State Psychiatric Institute, Psychiatry
Len Matin, Columbia University, Psychology
Andy Teich, Columbia University, Neurobiology
Tammy Moscrip, Columbia University, Psychology
John H. Hilton, Columbia University, Sergievsky Center
Jeevan Kalanithi, Rutgers University, CBMN
Sharon Himmanen, Lehman College
Karyl Swartz, Lehman College
Gina Arnone, Columbia University
Jennifer Mangels, Columbia University, Psychology
Sarah Shuwairi, Columbia University, Psychology
Jim Magnuson, Columbia University, Psychology
Emily Stern, Columbia University, Psychology
Domonick Wegesin, Columbia University, Sergievsky Center
Anja Soldan, Columbia University, Psychology
Jon Horvitz, Columbia University, Psychology
Johannes Schwaninger, Columbia University, Psychology
Mike Drew, Columbia University, Psychology
Peter Balsam, Barnard College, Psychology
Yaakov Stern, Columbia University, Sergievsky Center
Janet Metcalfe, Columbia University, Psychology
Brady Butterfield, Columbia University, Psychology
Chris Habeck, Columbia University, Sergievsky Center
Won Yung Choi, Columbia University, Psychology

Rapporteur: Kate Lynch

Dr. Cohen lectured on the neural processes involved in conflict monitoring and discussed the ways in which computational modeling and neuroimaging can be used to advance our understanding of cognitive control in particular and cognition in general. He began by making some broader comments on the importance of bridging the fields of cognitive psychology and neuroscience in our study of the mind/brain. He then presented some research done on cognitive control in his lab using computational modeling and neuroimaging in order to provide an example of how these two techniques have been used to study a particular phenomenon from an integrated cognitive neuroscientific approach.

Dr. Cohen explained that in order to bridge the fields of cognitive psychology and neuroscience, you must have a strong foundation in each of these respective fields. For over two decades cognitive psychologists have essentially rejected the notion that the brain had any relevance to the study of the mind. However, although this lack of consideration of the brain is striking, cognitive psychologists were able to explore the architecture of the mind and have established the foundation both in phenomena and methodical approaches to the study of behavior, both of which have served us well. On the other hand, the field of neuroscience proceeded with the idea that the mind could be understood by considering only the building blocks of the brain (i.e. functioning of calcium channels in a neuron), without concern for theories of how the mind works, developed from behavioral observation. However, a strong foundation in this field has been established.

Only now, after years of isolated research in cognitive psychology and neuroscience, has the bridge started to form between these two fields. The ability to use neuroimaging to study the function of an intact human brain has had a huge impact on the way in which research is done in the field of cognitive neuroscience, cognitive science, and neuroscience as well. Another tool, which has been given less emphasis but may be even more important, is computational modeling. Dr. Cohen argued that when dealing with a device as complex as the human brain, with nonlinear interactions between its parts, it is important to describe its functioning using equations, or by running simulations, to show that there is some empirical validity a particular idea of how the system is working. He went on to say that it is hubris to believe that we can understand and develop theories about the functioning of this device without the most powerful tools we have available: mathematics and computational modeling. Dr. Cohen suggested that the steadfast resistance of psychologists to the use of computational modeling and mathematics will ultimately be to our peril. Physicists have already begun to look into the issues of cognitive neuroscience using mathematics and computational modeling.

Dr. Cohen next provided an example of a specific research enterprise, the study of cognitive control, that takes advantage of neuroimaging and computational modeling in order to provide an explicit story of a piece of cognition that is interesting at the level of the mind, but that we come to better understand by considering it at the level of the brain.

Control is defined as the ability to behave in accord with internal states of intention or goals. For example, we have the ability to decide not to scratch a mosquito bite. We also have the capacity to adhere to that intention or goal and in so doing override another prepotent or reflexive response. Cohen recognizes that there is a much broader set of functions that one might put under the heading of control (i.e. carrying out sequential behavior, anticipating the future). He argued that the definition just given is, at the very least, an elemental component of a broader set of capacities. He argued that we must understand this specific type of cognitive control before we attempt to understand how we carry out more complex forms of control, such as that required when performing sequential behaviors.

Cognitive control is used to influence processing or performance in one part of the system. However, this cognitive control system is limited in capacity. In order for this system to function adaptively, there must be some part that monitors the processing space to determine where control is needed and then influences the control system such that its capacity is allocated judiciously. Despite the acknowledgment by Kahneman that a monitor is needed due to the limited capacity of the control system, up until the past 5-10 years there has been a scarcity of research investigating the mechanisms involved in monitoring, and their influence on control, at both the psychological and neurobiological levels. Data collected in Dr. Cohen's lab led him to consider the role of monitoring in the cognitive control system. He suggests that while on the one hand it makes sense to distinguish monitoring and control processes (as they involve control and evaluative/motivational processes respectively), on the other hand both must be considered together in order to understand how either one works. Cognitive control affects processing/performance, which is then monitored by a monitoring system. This monitoring system provides feedback to the control mechanisms. The nonlinear nature of these interactions is particularly conducive to the use of computational modeling.

Dr. Cohen showed a movie that exhibits how the combination of fMRI and EEG can be used to demonstrate the interactive, dynamic nature of brain activity. The image represented brain activity during the N-back working memory task. The activity associated with primary visual or motor processes was subtracted from the overall brain activity leaving only that activity associated with working memory or executive processes. Using both the information from ERP and fMRI Dr. Cohen was able to show that 100-250ms after the stimulus appears a burst of activity occurred in the right frontal areas. This frontal activity was followed by activity in the posterior parietal areas, which subsided right around the time that the subject responded. The complex of areas is involved in some kind of stimulus processing, evaluation, and the decision of how to respond. This data illustrated how dynamic the brain system is even in this simple working memory task. Computational modeling can be used to capture these complex dynamics. However, since models of complex phenomena involve many variables it can be difficult to test whether your model is right, especially when you only have a few dependant variables, and thus few constraints. There may be other accounts of the input-output relationship. The experimenter must run another study and manipulate the input some other way to see whether the output is still inline with predictions. This is a painful and time consuming process and so measuring internal processes using neuroimaging

tools can provide a much more direct and constraining test of the theory. While Dr. Cohen argues that mapping the terrain of the brain is a worthwhile endeavor, the main goal is to determine whether, and how, a process occurs in the brain, rather than where.

PART TWO:

The goal of the second part of the talk was to articulate in a formal way, in terms of mechanisms, how this elementary form of cognitive control may actually play out in the brain. Dr. Cohen provided an example of what is meant by control from the empirical literature and then presented a model that accounts for the data. He used that base to propose a mechanism that may be involved in monitoring and again used computational modeling and neuroimaging to provide support for the model. Dr. Cohen refers to this monitoring or feedback process as the “inner loop” of control. The inner loop of control adjusts the degree of control within the specified terms of the task. However, there is also an “outer loop” of control which monitors which of several tasks are worthwhile, when the task is completed, and at which point the next task should be started. Dr. Cohen hypothesized that the interaction between the locus coeruleus (a neuromodulatory nucleus in the brainstem) and the anterior cingulate may be the mechanism by which the “outer loop” functions.

In the Stroop task subjects are presented with color words in varying ink colors. The prepotent response is to read the word. However, subjects are capable of naming the color of ink the word is written in, if they are instructed to do so. The prepotent response to a Stroop stimulus is to read the word. However, when instructed, subjects are able to form an internal representation that indicates that the goal or intention is not to read the word. This ability can be taken as the canonical example of at least an elementary aspect of cognitive control. Dr. Cohen describes a real world example of a spatial Stroop effect in which a sign located on the upper left, indicates to the driver that he/she should proceed down the hill to the left. Phenomenologically, the spatial location interferes with the semantic content. Dr. Cohen next raised the question of how we can account for this type of phenomenon and the processing underlying it in terms how the brain processes information. He proposed the use of a “neural network” or “connectionist” model of neural processing that bears on behavior at the systems level. He argued that while there are some aspects of behavior that ultimately require a molecular level of analysis, the enterprise of neural network and connectionist modeling is to see how far we can get in the account of qualitative aspects in the brain, without committing ourselves to knowing every last detail.

Dr. Cohen presented a model for the Stroop task discussed earlier. He proposed two pathways that traverse several layers of representation or processing. In the first layer there are units that represent colors independently from orthographic form. These then map onto verbal responses by way of some sort of intermediate processing. An assumption is made that the mapping from orthographic form to verbal output was stronger and more consistent than the mapping of color onto verbal form. The enhanced strength is due to extensive experience using this pathway in real life tasks. Dr. Cohen

suggested that the phenomenology we experience specifically when the ink color conflicts with word meaning can be accounted for by the stronger weight assigned to this pathway and the mutual inhibition among representations in a given layer.

However, this model cannot systematically account for the ability to name the ink color, without some other source of input to the system. This other layer or module contains internal representations of the instructions or the task set, in this case the dimension of the stimulus that is relevant. Through the activation of the representation of the task set, by the prefrontal cortex (PFC), and by virtue of its connections to the relevant units in the processing pathways, priming occurs and results in faster flow of activity in the relevant pathway. This priming can outweigh the strength of the connections in the pathway representing the prepotent response. Representations of task set in the PFC bias processing, or the flow of activity, in the posterior structures, where the two pathways are represented. The role of the PFC, and the job of cognitive control, is to time posteriorly mediated processes in a way that avoids interference and leads to the execution of the proper set of processes for the task at hand.

The goal of the initial investigations was to provide a common mechanism that can explain several ideas at a psychological level relevant to attentional response selection in terms of the effects of task relevant information and competition from distracting or interfering information (i.e. inhibition). He explained that the PFC is not involved in the suppression of processing in a particular pathway, but rather supports, or primes the processing of information in the desired pathway such that the nodes in this pathway can compete more effectively with interference and guide activation along the pathways that are task relevant.

This model of cognitive control has done an excellent job at accounting for many phenomena in the literature pertaining to attention and control. The framework has led to some interesting predictions and has been useful in predicting and organizing findings in the neuropsychological literature regarding diseases that are known to be associated with the frontal and parietal cortices. However, at this point in the lecture Dr. Cohen shifted his focus to the topic of monitoring mechanisms. Dr. Cohen pointed out that when there is strong cognitive control the desired representation prevails at the output level and there is relatively low conflict. On the other hand when control is reduced, there is more conflict and the response is less reliable. A longer period of time is necessary to make a response and therefore more errors are made. The role of control in the model is to reduce conflict and produce the desired outcome.

The anterior cingulate is activated by competition among task relevant representations. That activation is actually correlated with errors and there are different paths for the association between the anterior cingulate activity, the error related negativity (ERN) and errors. In simple terms, activity in the anterior cingulate can be thought of as the product of the activity among competing units. For example, in the incongruent condition, when there is low control and the stimulus is the word green written in red ink, both the word form and the ink color are activated to some degree. When ink color and word form are equally active, the product of activity is maximized. The product of the activity in these

competing units is one of the things that drive anterior cingulate activity, and thus the anterior cingulate is most active when the units are equally active in the low control condition. When there is high control and less conflict the anterior cingulate is less active.

Botvinick and Cohen have shown that the model presented in this talk does an excellent job in accounting for the extant literature on the anterior cingulate activity and the ERN. They looked at all cognitive studies showing anterior cingulate activity for which a computational model already existed. These models were then taken and applied to the conditions used in Cohen's imaging studies to see whether the conflict measured predicted anterior cingulate activity. In all five cases, the conditions that were predicted to show increased anterior cingulate activity did in fact show an increase. In the cases where parametric data was available quantitatively accurate predictions about the degree of anterior cingulate activity were made.

The relationship between conflict and errors is instructive and leads to some counterintuitive predictions or explanations for some phenomena. Dr. Cohen reported that an imaging study done in his lab, using a task which manipulated conflict, showed not only enhanced anterior cingulate activity in error trials as compared to correct trials, but the same degree of activity in the high conflict and low conflict correct trials.

Dr. Cohen raised the question of what it is in the error trials that causes high conflict? Here is the explanation that the model provides. Assuming that errors are produced by noise at the response layer, every once in a while random noise will activate the incorrect node above threshold in the response layer. The initial, incorrect response is then elicited before the stimulus has been processed. After an erroneous response is given the stimulus comes on and the stimulus information is processed. This processing activates the correct response. Conflict arises due to the competing activation from the random noise and actual stimulus processing. However, the conflict should occur after the response because the response occurs before the stimulus information is acquired. The ERN occurs due to the post response processing of the stimulus in an error trial. The idea that it is conflict, rather than the error itself which leads to anterior cingulate activity and the ERN, leads to a whole group of predictions, many of which have been tested in the literature.

There was resistance in the field to this idea and several studies attempted to show results that contradict this theory, and test some of the predictions. Some phenomena were reported contradict the model. Dr. Cohen focused on one particular finding showing that a larger ERN to the congruent than to the incongruent stimuli in error trials. Dr. Cohen originally predicted a larger ERN for incongruent than congruent error trials, because greater conflict was predicted in this condition. Scheffers and Coles ran an experiment to investigate these predictions and found the exact opposite. Dr. Cohen highlights the importance of modeling using this particular situation. In this particular case Dr. Cohen's intuition was incorrect, and when we ran the model on the data, the model predicted the same result that they found experimentally: greater conflict and thus greater anterior cingulate activity in the congruent than the incongruent error trials.

Upon closer analysis of the model the reasons for this effect became clear. In an error trial a preburst of activity (noise) occurs on the incorrect channel (inconsistent with the task set). During an incongruent trial additional activity occurs in the incorrect channel however, there is little benefit from this increased activation above and beyond that already produced by the noise. In the congruent condition, there is no additional activation in the incorrect pathway. Only the correct pathway is activated. The activation of the incorrect pathway by noise and the activation of the correct pathway by the congruent stimulus produced a greater degree of conflict than in the incongruent condition. Thus the model predicted exactly what was observed in the data.

On a correct trial greater conflict is predicted in the incongruent trials than the congruent trials. The response on a correct trial is stimulus driven and thus conflict results when the stimulus word interferes with color naming in the incongruent condition. Dr. Cohen reported that a negativity for correct incompatible trials (N2C) was found in the literature. This negativity occurred before rather than after the response as would be predicted by the model.

Dr. Cohen next discussed some predictions that arose from the model at this point. The representations of task sets, goals, and intentions, are located in the PFC. These representations modulate activity in other parts of the system. The AC, on the other hand, may be monitoring the outcome of processing to determine whether or not there is conflict. Dr. Cohen emphasized that if the PFC and AC perform separate functions, one should be able to dissociate the activity of these two areas, under the right conditions. Cohen ran an experiment in which subjects were given an instruction at the beginning of each trial (either name color or read word). There was a delay and then the typical congruent or incongruent Stroop stimulus was shown.

Dr. Cohen predicted engagement of the PFC in response to these instructions and more specifically greater activity to the color naming instruction than the word reading instruction, because word reading is the default response. No differential activity was predicted in the anterior cingulate to the instruction stimulus. Conversely, Dr. Cohen predicted that when the Stroop stimulus appeared the AC would be more active in response to the incongruent Stroop stimulus than the congruent one. This result has been reported several times in the literature. No differential activity was predicted in the PFC to the congruent vs. incongruent Stroop stimuli.

As predicted, during the instruction and delay periods, the PFC showed greater activation in the color naming than in the word-reading task, and no differential activity was found in the AC. Dr. Cohen hypothesized that once instructed subjects prepare and prime a response set, a subset of possible responses (red, green, etc.) There is some behavioral evidence that subjects do this. Dr. Cohen went on to hypothesize that activation of the nodes representing all possible responses, may cause competition between the elements of the response set and thus AC activity. During the stimulus response period, as predicted, greater activity was found in the AC for the incongruent than the congruent stimulus. The pattern of cingulate activity varies as a function of conflict. PFC activity always occurs when there is AC activity. However, in this task there was no differential

activity in the PFC between congruent and incongruent trials. A true double dissociation was found between the two structures in accord with theoretical predictions.

In summary, Dr. Cohen has provided neuroimaging evidence that the PFC is involved in the allocation of control and anterior cingulate is involved in conflict monitoring. Some counterintuitive findings were found concerning this part of the monitoring mechanism. Dr. Cohen looked at the effects of adjustment the conflict monitor on the fidelity of the prefrontal representations. As was done with the imaging data, members of Dr. Cohen's lab used the basic model to simulate the effect of changes in the conflict measure on changes in activity in the prefrontal layer depending on the time course over which you integrate the conflict. If only conflict on the last trial is used adjust the PFC layer you do not get a good fit to the data. However, if you time average using exponential decay, you get an excellent fit.

The model makes some specific predictions in the imaging domain about trial to trial interactions, the way in which conflict on one trial affects reduction in PFC activity on the next trial. Dr. Cohen is now using fMRI to determine the trajectory of activity over many trials.

Dr. Cohen pointed out that the AC is a very large structure and that it is likely involved in broader functions than just conflict monitoring. Imaging studies have produced very consistent results, but these results are limited in scope. Dr. Cohen posed the question of what, if anything, the AC is doing as a whole. If there is something else, how does conflict monitoring fit within that broader function. Dr. Cohen suggested that perhaps conflict monitoring is one component of a function that is responsible for monitoring internal states for signs of threat in much the same way the amygdala is involved in monitoring the external world for signs of threat. Many studies have found AC activity in response to pain stimuli and error feedback, both of which are signs of threat. Conflict monitoring may act as an early warning signal, so that outcomes can be anticipated.

In the classic model of depression if you take away control from your subjects they get depressed. It is not clear how good a model that is for all forms of depression. Certainly the loss of a family member would lead to the acute sense that you have lost control. However, many cases of depression develop over time, and are not necessarily in response to a specific event. Dr. Cohen suggested that the second type of depression may result in part from a form of endogenous learned helplessness resulting from a disconnection between the monitor and the executive control. A loss of control would result from the inability to adjust control even when the monitor indicates that the system is performing poorly. This disconnection could lead to the same sense of loss of control or learned helplessness that external loss of control produces. Dr. Cohen has begun to investigate this idea that depression may have to do with the disconnection between the AC and PFC.

In conclusion, Dr. Cohen showed that this explicit modeling framework can be used to describe the self-regulating mechanisms for cognitive control that integrate executive and evaluative functions. Initial neurobiological findings have shown that the PFC is

involved in allocation of control and biases activity along task-relevant pathways. The AC is involved in monitoring conflict to determine when adjustments in control are needed. Although Dr. Cohen did not mention this in his talk, the adjustments are made by way of neuromodulatory regulation in the locus coeruleus. Lastly, Dr. Cohen has provided an example of the way in which the combined approach of computational modeling and imaging can be used to generate and test novel hypotheses in general (e.g. the function of the ACC, and the LC) and in particular identifying the relationship between neural substrates and cognitive function.

Questions:

Moeller referred to a book from 1960, by Miller, Galanter and Prebram which discussed the idea of feedback hierarchies and cognitive control and asked Cohen to expand on the relevance of the issues in this book to the discussion today.

Cohen responded that his historical introduction may have been hyperbolic, and that indeed the book mentioned by Moeller is very relevant to the discussion. He continued by saying that modern cognitive science has not heeded the notion of control very much, and that there is a lot to be gained by considering these issues. Cohen explained that members of his team are very well versed in the history of control theory, and it's modern form in nonlinear dynamical systems analysis. The goal is to understand how these processes work and try to develop even more sophisticated predictions that would enhance confidence in specific models and allow us to adapt these models in ways that would better fit behavior.

Magnuson asked how high up in the system control structures will be needed.

Cohen responded that the relationship described so far between conflict and control could be thought of as monotonic. However, one can imagine situations in which conflict persists for so long that reward value may decrease, more interesting awards come along, and thus in longer makes sense to keep trying. In this case a reduction of control would be beneficial.

A couple of observations led Cohen to believe that the locus coeruleus may be playing a role in this function. First of all, the locus coeruleus seems to have two phases in the awake animal, an intermediate phase of tonic activity in which you see phasic bursts of firing in response to a stimulus and really good task performance; and another phase in which the tonic activity is higher, without phasic bursts. This second phase seems to correlate very tightly with more variable motor performance in the animal, with false alarms and with the amount of scanning or eye movements.

This finding led Cohen and his colleagues to believe that the LC is mediating between focused and variable, scanning states. As this work began to mature it occurred to Cohen and his colleagues that this function may represent an adaptive system at the outer loop level. Cohen proposed that while the AC is averaging conflict over two time points, a short time frame and a long time frame. In the short time frame if you have a small burst

of conflict in the context of very little overall conflict and if the task is still rewarding, the subject is likely to try harder.

Over a long time frame conflict accrues, and the subject is more likely to stop trying. The AC is integrating conflict over two frames, and this information may indicate to the LC whether to move into this more focused state, or to broaden the scope of attention and give the system a chance to leave the current task set behind and try a new one. In line with the predictions of this theory, primate studies have shown that there is a huge projection from the AC to the LC. Cohen pointed out that in fact, the largest cortical projections to the LC are from the AC. The theory would also predict that changes in LC state would occur as a function of AC activity. This result has not yet been found, however, Cohen is currently investigating these predictions using physiological recording studies.

Terrace commented that the reaction time (RT) data was only presented in once specific study. Terrace asked Cohen to comment on the RT data in the other work presented during the lecture today.

Cohen replied that here were predictions about the RT data in every simulation.

Terrace asked whether it is usually the case that RTs on error trials are considerably larger than on correct trials.

Cohen replied that not only was the mean RT larger on error trials, but the distribution of RTs varied as well. Cohen explained that his research team is careful to look at the behavioral data before simulating any task.

Fossella asked about the overall capacity of the system.

Cohen replied that capacity is a construct that has been valued quite a bit in the cognitive literature and by neuroscientists. However, Cohen said that he believes this concept is difficult to deal with. It is difficult to operationalize capacity. Cohen proposed that his model may operationalize difficulty, with difficulty defined by the degree of conflict produced by a task, given the processing systems capability of performance. However, Cohen said that measuring capacity is difficult. Perhaps capacity can be operationalized as the sustainability of the representation in the PFC needed to establish the task set.

Perhaps capacity is defined as the number of competing representations that can be sustained at the same time in a given part of a system. Then we can begin to ask what damage to individual of units or connections does to this capacity. Cohen emphasized that it is not obvious that as you delete units, you reduce capacity. It is possible that gain capacity increases. Cohen concluded his response by expressing that he is uncomfortable discussing this issue without having some mechanistic quantitative account of what it means.

Davidson asked about the flexibility of the system.

Cohen responded that this question brings up the role of dopamine. Within a task there are goals, and subgoals and there must be a mechanism to update the system, such that when a new representation is active, the system can determine how it relates to the current representation. Dr. Cohen suggested that the basal ganglia-PFC loop may be involved in updating the system with information about subgoals. Cohen said that the flexibility appears to lie in this BG-PFC system.