Implicit Learning Deficits in Autism: A Neurocomputational Account

Trent Kriete & David C. Noelle (tkriete,dnoelle@ucmerced.edu)
Program in Cognitive Science, University of California, Merced, 5200 N. Lake Road
Merced, CA 95343 USA

Abstract

Experiments using the Serial Response Time Task (SRTT) have suggested that implicit learning is impaired in people with autism. Implicit learning is learning that occurs without explicit awareness of the knowledge being acquired. Researchers have suggested that poor implicit learning could be a major factor in other profound learning difficulties exhibited by people with autism. In this report, we use a neurocomputational model of the SRTT to show that disturbed interactions between the mesolimbic dopamine (DA) system and the prefrontal cortex (PFC) may underlie the implicit learning problems observed in autism. This model is shown to fit reaction time data from the literature for both individuals with autism and normally developing controls. This work expands on a previous body of research showing that abnormal DA/PFC interactions can account for a disparate collection of autistic behavioral patterns, suggesting that a common neurological mechanism might produce executive dysfunction, stimulus overselectivity, and impaired implicit learning in this population.

Keywords: Autism; Implicit Learning; Serial Response Time Task (SRTT); Prefrontal Cortex (PFC); Dopamine (DA)

Introduction

Autism is a complex developmental disorder characterized by deficits across physical, social, and cognitive domains. Cognitive difficulties are found in tasks assessing executive function, "mind reading" abilities, integration of information, attention, and generalizing learned abilities to novel contexts (Hughes, Russell, & Robbins, 1994; Baron-Cohen, Leslie, & Frith, 1985; Frith, 1989). In addition, physical motor abnormalities, an increased prevalence of seizure disorders, motor stereotypies, and repetitive behaviors often accompany the diagnosis (Rinehart, Bradshaw, Brereton, & Tonge, 2001; Tuchman & Rapin, 2002). Complicating matters further, people with autism often possess "islets" of spared cognitive functioning in some areas, including visiospatial tasks such as the embedded figures test and Weschler's block design task (Shah & Frith, 1983). Indeed, the diversity and variance of the traits of autism has led to a recent proposal to abandon any attempt to find a single "monolithic" cause underlying all aspects of the disorder (Happe, Ronald, & Plomin, 2006).

While we are sympathetic to the notion that *all* behavioral aspects of autism are unlikely to have a single common neurological cause, our previous explorations of computational models of the prefrontal cortex (PFC), and its interactions with the mesolimbic dopamine (DA) system, have suggested that a disparate collection of behavioral patterns observed in autism might stem from a common deficit in these neural systems. We propose that abnormal PFC/DA interactions may underlie multiple aspects of this disorder, providing a neuroscientific account of how these aspects are related. Our previous computational work has demonstrated how deficits in PFC/DA interactions can account for patterns of *executive dysfunction* in autism (Kriete & Noelle, 2005, 2006) and how the same neural deficits can explain *stimulus overselectivity* in people with autism along with the associated difficulties in generalizing learned skills to novel contexts (Kriete & Noelle, 2008). In this paper, we present a neurocomputational model of implicit learning performance in people with autism, demonstrating that abnormal PFC/DA interactions can explain the observations in this domain, as well.

Implicit learning is learning that occurs without any awareness of the specific knowledge acquired during the process. Researchers have suggested that people with autism have a core deficit in their ability to implicitly learn about the inherent relationships that exist between objects and situations in the world (Mostofsky, Goldberg, & Landa, 2000; Klinger, Klinger, & Pohlig, 2006). Klinger et al. argue that impaired implicit learning results in difficulties in recognizing the relationships that exist across experiences, likely leading to problems forming general knowledge about categories of items and types of situations. Difficulties in generalizing learned knowledge to new situations are commonly observed in people with autism, and these difficulties frequently act as a central obstacle to the development of behaviors needed for autonomy and independent living. Thus, a precise characterization of the mechanisms responsible for these generalization deficits would be very valuable to any effort to design ways to mitigate these serious issues in people with autism.

In this paper, we first describe our general computational account of PFC/DA interactions, and we review our previous work, showing that abnormalities in these neural systems can account for the patterns of executive dysfunction and stimulus overselectivity observed in autism. We then review a common paradigm for assessing implicit learning, the Serial Response Time Task (SRTT), and we revisit experimental results that make use of this task to demonstrate impairments in implicit learning in people with autism. A neurocomputational model of healthy performance on the SRTT is then described, and the results of introducing a PFC/DA deficit into this model are reported. The model results are fit to both healthy and autistic SRTT performance data appearing in the literature, demonstrating the ability of a PFC/DA interaction abnormality to account for the lack of implicit learning in autism. We close with a general discussion.

Previous Work

Our proposal is that deficits in the mesolimbic dopamine (DA) system and/or abnormalities in how dopamine modulates the prefrontal cortex (PFC) can account for the behavioral profiles of people with autism across a diverse range of

task domains. Thus, we suggest that some behavioral deficits that have previously been seen as stemming from separate psychological mechanisms might have a common neurological cause. Our strategy for demonstrating the feasibility of this proposal is to show that a broad array of state-of-the-art neurocomputational models of healthy human performance all exhibit autistic patterns of responding when they are minimally modified to reflect a failure of DA to properly modulate PFC. In this section, we review our general theory of the functional role played by PFC/DA interactions, as well as our previous modeling demonstrations involving *executive dysfunction* and *stimulus overselectivity* in autism.

Interactions Between DA & PFC

We share an account of PFC function that was initially developed primarily to explain the role of PFC in cognitive control and cognitive flexibility (Cohen, Dunbar, & McClelland, 1990; Braver & Cohen, 2000; Rougier, Noelle, Braver, Cohen, & O'Reilly, 2005). Cognitive control is the ability to guide behavior according to explicit goals or rules, especially when doing so is in conflict with more automatic or prepotent tendencies. Cognitive flexibility describes the ability to appropriately adapt cognitive control in response to shifting task contingencies. The PFC has been broadly implicated in cognitive control and cognitive flexibility. In our models, the PFC supports cognitive control by actively maintaining abstract rule-like representations that provide top-down modulation of more posterior brain areas, modifying the regular behavior of these posterior pathways so as to overcome their usual automatic patterns of responding (Cohen et al., 1990). Biologically, the active maintenance of frontal control representations is supported by dense patterns of recurrent excitation in the PFC, as well as intrinsic maintenance currents. Computational models of these neural circuits have shown that the active maintenance of control representations and the flexible adaptation of control are at odds, with the mechanisms that maintain PFC representations, and protect them from distracting inputs, acting as an obstacle to the rapid updating of PFC contents in response to shifting contingencies. Thus, in order to achieve cognitive flexibility, a separate mechanism is needed to rapidly update the actively maintained PFC control state in a task-appropriate manner.

A useful analogy for this flexible updating mechanism is that of a "gate" in a fenced enclosure. When cognitive control must be strong, the gate is closed, keeping out distracting inputs that might compromise the needed PFC control signals. When the current control state is no longer appropriate, the gate opens, allowing the old control state to escape and allowing a new control representation to enter the PFC via its inputs. In order for the PFC to maintain situation-appropriate control as contingencies change, a neural mechanism is needed that can learn to adaptively open and close the gate on PFC in a task-appropriate manner. Some researchers have suggested that the mesolimbic dopamine system may play a central role in learning to control this gate (Braver & Cohen, 2000; O'Reilly & Frank, 2006). Dopamine cells have been found to carry reward prediction information critical for learning associations between behaviors and reward (Montague, Dayan, & Sejnowski, 1996), and the DA projections to PFC have been viewed as a likely neural implementation of the gating signal needed to flexibly adjust the control state of PFC (Braver & Cohen, 2000).

Under this account, DA interactions with PFC drive the flexible updating of control. Inflexibility arises when these interactions are disturbed, frequently resulting in PFC perseverating on control representations that are no longer appropriate. This insight, along with evidence of DA abnormalities in autism, has led us to investigate the degree to which the perturbation of DA/PFC interactions naturally leads to patterns of behavior observed in people with autism. Our previous computational modeling work has shown that this mechanism is sufficient to explain various aspects of *executive dysfunction* and *stimulus overselectivity* in autism (Kriete & Noelle, 2005, 2006, 2008).

Executive Dysfunction

People with autism are impaired across a broad range of cognitive tasks that have been associated with executive control processes. Indeed, the *Executive Dysfunction (ED)* theory of autism seeks to explain many of the behavioral patterns exhibited by these individuals in terms of a failure of executive control over behavior (Hughes et al., 1994).

There is extensive evidence that the prefrontal cortex plays an important role in executive control. Along with the central claim of ED, this suggests that the root cause of many autistic behavioral patterns may lie in abnormalities in this region of the brain. This view of ED suggests that the irregular development of prefrontal cortex may underlie the patterns of cognitive performance seen in autism.

A more detailed examination of autistic behavior reveals that not all forms of executive processing are impaired, however. A perplexing aspect of the executive profile demonstrated by people with autism is that cognitive flexibility has been shown to be impaired while fundamental cognitive control remains robust and relatively unaffected. A classic measure of cognitive control is the Stroop task (Stroop, 1935), and a common measure of cognitive flexibility is performance on the Wisconsin Card Sort Test (WCST) (Berg, 1948). Persons with autism have been shown to exhibit poor WCST performance, but they exhibit no more interference on the Stroop task than healthy controls (Ozonoff & Jensen, 1999). This dichotomy challenges the notion that autistic behavior is the result of a global impairment of executive processes, perhaps mediated by frontal abnormalities.

One clear way to resolve this issue is to posit separate mechanism for cognitive control and the flexible adaptation of control. In people with autism, the ability to actively maintain information and influence behavior, the PFC, is intact. However, the ability to flexibly adapt control, mediated via the DA system, is impaired.

In order to demonstrate the viability of this account of executive dysfunction in autism, we made use of an exist-

ing neurocomputational model of PFC/DA interactions that had been shown to fit the performance of both healthy and frontally damaged humans on both the Stroop and WCST tasks (Rougier et al., 2005). Simply weakening the influence of DA on PFC in this computational model was sufficient to both qualitatively and quantitatively capture autistic performance on both Stroop and WCST (Kriete & Noelle, 2005, 2006). More specifically, reducing the gating effect of the DA signal selectively impaired the ability of the system to update the contents of the PFC, which is vital for the WCST. However, there was no effect on the ability of the PFC to influence more posterior brain areas and assist overcoming prepotent responses during the Stroop task. This computational modeling result suggested that executive deficits in autism may be mediated by weakened PFC/DA interactions.

Stimulus Overselectivity

Stimulus overselectivity is said to occur when an overly restricted set of features of the environment come to drive behavior. While some associations between environmental features and appropriate action are acquired, many reliable cues to action are not learned. This phenomenon was first documented in the early 1970s in people with autism (Lovaas, Schreibman, Koegel, & Rehm, 1971). Overselective behavior in people with autism has been seen as a plausible explanation for the problems many individuals with autism exhibit when they are expected to generalize learned behaviors to novel situations. In such situations, restricted, often irrelevant, portions of the environment become tightly coupled with the performance of the desired behavior. If this restricted portion of the environment is not consistently available to the individual, generalization to new settings will suffer. For example, an individual with autism might successfully learn to order a meal from a simulated cashier in a laboratory but fail to generalize the required actions to situations that fail to possess some key feature, like a particular item of furniture or idiosyncratic word usage on the part of the cashier. This inability to appropriately generalize learned skills is a major focus of many behavioral intervention techniques.

In Kriete & Noelle (2008) we presented a computational account of how overselective representations can develop in more posterior brain areas when the PFC is unable to flexibly update its contents, due to inappropriate DA modulation. We used a relatively simple Leabra (O'Reilly & Munakata, 2000) neural network model of conditioning that included a modulation of the stimulus-response mapping based on the contents of a layer of simulated PFC neurons. In this case, the control signals actively maintained in PFC encoded executive attention, selectively highlighting specific aspects of the current stimulus. With healthy PFC/DA interactions, the PFC contents flexibly switched from one aspect of the current stimulus to another, allowing the network to learn about all relevant aspects of the stimulus. When the PFC was unable to flexibly and appropriately update its contents, however, representations in cortical areas downstream from the PFC developed so as to be dominated by an overly restricted, or possibly even irrelevant, subset of features present in the stimulus. Poor generalization was exhibited by these networks, due to these abnormal cortical representations. In the model, the inability to flexibly update PFC increased the likelihood that the only environmental associations learned involved spurious correlations (i.e., idiosyncratic features of the training environment), with other, more broadly relevant, features escaping the attention of the network. Subsequent dependence on such spurious correlations crippled generalization performance. The model's behavior was favorably compared to previously reported laboratory data involving people with autism and healthy controls. This finding provides an additional example of how a common neurological deficit, dysfunctional PFC/DA interactions, can help to bridge theoretical gaps across behavioral domains within autism research.

The Serial Response Time Task (SRTT)

In addition to executive dysfunction and stimulus overselectivity, is it plausible that abnormal PFC/DA interactions can also account for the deficits in implicit learning observed in people with autism? We address this question by investigating the effects of PFC/DA dysfunction on the performance of a psychological test commonly used to assess implicit learning abilities: the Serial Response Time Task (SRTT).

In a common version of this task, participants are presented with four buttons, with exactly one button illuminated at any one time. Participants are asked to simply press the currently illuminated button as quickly and accurately as possible. Once a button is depressed, a new button is illuminated, prompting the participant to press the new button, and this sequence of cued button presses continues for a block of 80 responses, with an experimental session consisting of five of these blocks. The illumination order of the buttons is the key manipulation of the SRTT. During the first and the final (fifth) block the order in which the buttons are illuminated is random. However, during blocks 2, 3, and 4 there is a hidden pattern in the responses that are required. This hidden structure is apparently detected by many healthy participants, as there is a significant reduction in the reaction time required to press the correct button across blocks 2, 3, and 4. Importantly, this reduction in reaction time does not occur during the randomized first and fifth blocks. The common interpretation of these results is that learned knowledge of the hidden sequential pattern allows participants to better "anticipate" which button will be illuminated next, allowing them to prepare this upcoming action and, thereby, speed their response. Knowledge of the hidden structure is seen as "implicit", however, as most participants claim no explicit knowledge of the sequential pattern (Cleeremans & McClelland, 1991).

People with autism, however, do not show marked improvement during the intermediate blocks of the SRTT, providing support to the claim that autism impairs implicit learning abilities (Mostofsky et al., 2000). While this result is interesting in its own right, we do not yet have an understanding of the biological mechanism(s) behind this deficit.

Some insight might be gained from the neuropsychological literature involving the SRTT. Specifically, deficits in tasks assessing implicit learning have been linked to damage to the cerebellum. This is intriguing, as there is ample evidence of cerebellar abnormalities in people with autism (Courchesne et al., 1994). However, other tasks traditionally associated with the cerebellum, such as judgment of timing, show no differences between people with autism and normally developing controls (Mostofsky et al., 2000). Recently, evidence has emerged suggesting that PFC and the basal ganglia may be important players in implicit learning as well (Matsumoto, Hanakawa, Maki, Graybiel, & Kimura, 1999; Pascual-Leone, Wassermann, Grafman, & Hallett, 2004). It is this latter connection that we will pursue, here, using an established computational model of the SRTT to investigate the possibility that PFC/DA abnormalities may give rise to the implicit learning problems observed in people with autism.

Modeling Implicit Learning

Seminal work on modeling healthy SRTT performance has been conducted by Cleeremans et al. (1991). In these neurocomputational models, a simulated neural circuit is presented with an input that encodes the currently illuminated button, and the output of this circuit is read as the system's expectation for the next button to be illuminated. This input to output mapping is mediated by a collection of hidden units, and synaptic learning methods are used to improve this mapping with experience. Importantly, these networks also include a "context layer" of neural units which can learn to actively maintain information about the history of previously presented inputs, allowing the model to base its predictions on more than the currently illuminated button. The activation levels of the neural units in the context layer are set to be "copies" of the hidden unit activation levels whenever a new input is presented, making the Cleeremans model of SRTT performance essentially a simple recurrent network (SRN) (Elman, 1990) trained to predict the next button press. The schematic network architecture is shown in Figure 1. This model has provided good fits to healthy human performance on the SRTT (Cleeremans & McClelland, 1991).

Since the sequential structure in the intermediate blocks of the SRTT is often complex, the information provided by the context layer is vital for the success of the model. Importantly, the context layer in this model plays an identical functional role to the PFC in other models, actively maintaining information that can be used to modulate an input-output mapping. In our previous executive dysfunction model, the PFC maintained information about the currently relevant stimulus dimension (e.g., "focus on the ink color" in Stroop or "sort cards based on shape" in WCST), so as to modulate performance. In our stimulus overselectivity model, the PFC maintained information about the current stimulus feature to receive executive attention, modulating how the stimulus was processed. In this model, the context layer maintains information about the preceding button presses, allowing that information to modulate the prediction of the next button. The main difference between our previous PFC models and this SRTT model involves the timing with which the contents are updated. In our previous models, the PFC was updated in a dynamic fashion, based on learned task contingencies. In this model, the context layer is updated with each new input presentation. Thus, the SRN context layer is analogous to the PFC in our previous models, with the updating "gate" forced to open with each new input, as previously suggested by other researchers (O'Reilly & Frank, 2006).

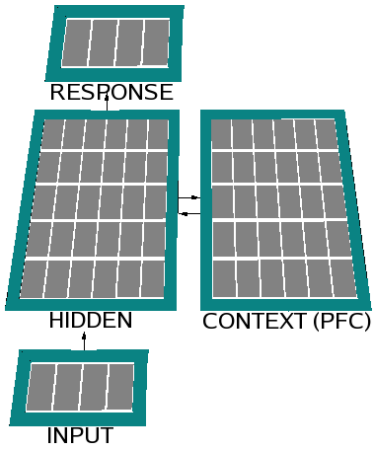

Figure 1: Network Diagram of SRTT Model

It is important to note that, in order to capture the relevant sequential information, the SRN must update the context layer in a fast and appropriate manner. This flexible updating of contextual information is precisely the cognitive mechanism we hypothesize to be suspect in people with autism. By restricting the ability of the SRN to update the context layer, mirroring the PFC updating failures that arise with weakened PFC/DA interactions in our other models, we expect to capture the performance of people with autism.

We made small modifications to a previous implementation of the Cleermans SRTT model which uses the biologically grounded Leabra framework, reducing the original implementation's 10-unit inputs and outputs to only 4, to capture the structure of the 4-button SRTT (O'Reilly & Munakata, 2000). The resulting network is shown in Figure 1. In this model, an Input Layer represents the four distinct buttons. The Hidden Layer learns the prediction mapping and provides a modeled abstraction of posterior brain systems. A Response Layer encodes the prediction output of the network. Additionally, a Context Layer provides a top-down influence on processing within the Hidden Layer, reflecting role of PFC.

In order to model the performance of people with autism, we restricted the probability of successfully updating the Context Layer (PFC) upon each input presentation. Normally, the Context Layer is updated with each input, but our autism model only updated the layer with some fixed probability which was less than one. This is analogous to reducing

the efficacy of the DA-based gating signal to the PFC. Restricting the updating of the PFC in this manner makes the temporally extended information normally contained within this layer much less reliable, making the learning of complex sequential structures much more difficult for the network.

The measure of interest in the SRTT is the response time of the participants throughout the training blocks. In order to compare model performance to human reaction times, Cleeremans et al. (1991) translated network outputs into a probability distribution over the four buttons using a Luce choice ratio (Luce, 1963) and then linearly scaled the error between this prediction distribution and the actual outcome (i.e., the next button illuminated) to produce a modeled response time. This procedure assumes that there is a linear increase in response time with prediction error. We used this method, as well, introducing three free parameters for fitting the model to data: a linear scaling parameter from error to milliseconds, a base response time (when error is zero) for the healthy model, and a base response time for the autism model. Note that different base response times were used for the normally developing and autism cases in order to capture the difficulty people with autism regularly exhibit when initiating movements (Rinehart et al., 2001).

Results

Network simulations were repeated 100 times in each of the experimental conditions, with initial synaptic weights randomized for each repetition. Average performance results for each block were compared to previously reported response time data for both people with autism and normally developing controls (Mostofsky et al., 2000). A grid search was performed over the possible probabilities of updating the Context Layer, testing probabilities from 0.0 to 1.0 in steps of 0.1, with the three linear scaling parameters optimized to reduce sum-squared deviation from the human response time data. The updating probability, and associated scaling parameters, that produced the lowest sum-squared deviation from the human data was identified as the best fit model.

The simulation results match human performance both qualitatively and quantitatively, providing evidence that impairments in PFC updating can result in implicit learning deficits like those seen in people with autism. When the healthy network is restricted to perfectly update its Context Layer (i.e., with probability one), the corresponding best-fit probability for the autism network is 0.5, with an SSE of 652.[1] The corresponding scaling parameter from error to reaction time is 261.4, the healthy base time is 458.6 msec, and the autism base time is 534.5 msec. The resulting modeled reaction times, along with human data, is shown in Figure 2.

[1] If the healthy network is allowed to update imperfectly, as well as the autism network, the best fit arises when the healthy network updates with 0.6 probability and the autism network updates with 0.2 probability, producing an SSE of 549. Unfortunately, since the variance of the human data was not reported in Mostofsky et al. (2000), we cannot assess if these parameters are reliably better at fitting the human data than the 1.0/0.5 case.

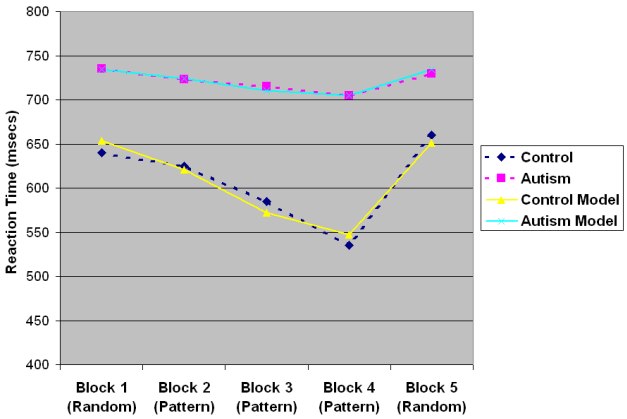

Figure 2: Scaled Model Results & Human Behavioral Data from Mostofsky et al. (2000)

A repeated measures ANOVA was conducted on blocks 1, 2, 3, and 4 of the model results, and a significant Group by Block interaction ($F(3,97) = 62.007$; $p < 0.000001$) was detected. From these results we can conclude that the networks simulating autistic performance demonstrated significantly less learning over the crucial training blocks of 2, 3, and 4, as compared to the networks allowed to properly update their PFC-like Context Layers. Thus, clear implicit learning deficits were present in the autism model.

Discussion & Conclusion

The breadth of behavioral and neurological abnormalities discovered in people with autism is almost staggering. Using computational models, constrained by our existing knowledge of biology, is a relatively untapped resource for exploring the neurological underpinnings of autism. Utilizing these tools in our investigations has provided us with insights into how a single neurological mechanism — the improper updating of the PFC caused by disrupted interactions with the mesolimbic DA system — can account for behavior across a variety of previously unrelated domains of autistic behavior.

The modeling results presented in this report suggest that, in people with autism, implicit learning deficits may be driven largely by abnormalities in DA/PFC interactions, causing inflexibility in the updating of contextual information. Without the proper updating of contextual information, it is essentially impossible to properly integrate temporally separated pieces of information, such as the order of items in a sequence. Thus, our computational account highlights how PFC/DA dysfunction can lead to problems with information integration. This is particularly interesting, since one prominent behavioral theory of autism, *Weak Central Coherence*, posits that deficits in integrating contextual information lay at the core of this disorder (Happe, 1999). It is a major point of our future research to investigate whether abnormal PFC/DA interactions can account for the various other behavioral patterns currently cited as evidence of weak central coherence.

Acknowledgments

The authors extend thanks to the members of the Computational Cognitive Neuroscience Laboratory at the University of California, Merced at to three anonymous reviewers.

References

Baron-Cohen, S., Leslie, A. M., & Frith, U. (1985). Does the autistic child have a theory of mind. *Cognition*, *21*, 37–46.

Berg, E. A. (1948). A simple objective test for measuring flexibility in thinking. *Journal of General Psychology*, *39*, 15–22.

Braver, T. S., & Cohen, J. D. (2000). On the control of control: The role of dopamine in regulating prefrontal function and working memory. In S. Monsell & J. Driver (Eds.), *Control of cognitive processes: Attention and performance XVIII* (pp. 713–737). Cambridge, MA: MIT Press.

Cleeremans, A., & McClelland, J. L. (1991). Learning the structure of event sequences. *Journal of Experimental Psychology*, *120*(3), 235–253.

Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing model of the stroop effect. *Psychological Review*, *97*(3), 332–361.

Courchesne, E., Townsend, J., Akshoomoff, N. A., Saitoh, O., Yeung-Courchesne, R., Lincoln, A., et al. (1994). Impairment in shifting attention in autistic and cerebellar patients. *Behavioral Neuroscience*, *108*(1), 848–865.

Elman, J. (1990). Finding structure in time. *Cognitive Science*, *14*, 179–211.

Frith, U. (1989). *Autism: Explaining the enigma*. Oxford: Blackwell.

Happe, F. (1999). Autism: Cognitive deficit or cognitive style? *Trends in Cognitive Sciences*, *3*(6), 216–222.

Happe, F., Ronald, A., & Plomin, R. (2006). Time to give up on a single explanation for autism. *Nature Neuroscience*, *9*(10), 1218–1220.

Hughes, C., Russell, J., & Robbins, T. W. (1994). Evidence for executive dysfunction in autism. *Neuropsychologia*, *32*(4), 477–492.

Klinger, L. G., Klinger, M. R., & Pohlig, R. A. (2006). Implicit learning impairments in autism spectrum disorders: Implications for treatment. In J. M. Perez, P. M. Gonzalez, M. L. Comi, & C. Nieto (Eds.), *New developments in autism.* London: Jessica Kinglsey.

Kriete, T., & Noelle, D. C. (2005). Impaired cognitiv flexibility and intact cognitive control in autism: A computational cognitive neuroscience approach. In *Proceedings of the 27th annual conference of the cognitive science society* (pp. 1190–1195). Lawrence Erlbaum.

Kriete, T., & Noelle, D. C. (2006). Dopamine and the development of executive dysfunction in autism. In *Proceedings of the 5th international conference on development and learning.* Lawrence Erlbaum.

Kriete, T., & Noelle, D. C. (2008). Modeling the development of overselectivity in austism. In *Proceedings of the 7th international conference on development and learning.* Lawrence Erlbaum.

Lovaas, O., Schreibman, L., Koegel, R., & Rehm, R. (1971). Selective responding by autistic children to multiple sensory input. *Journal of Abnormal Psychology*, *77*(3), 211–222.

Luce, R. D. (1963). Detection and recognition. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.), (pp. 103–189). New York: Wiley.

Matsumoto, N., Hanakawa, T., Maki, S., Graybiel, A. M., & Kimura, M. (1999, August 1). Nigrostriatal dopamine system in learning to perform sequential motor tasks in a predictive manner. *Journal of Neurophysiology*, *82*(2), 978–998.

Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive hebbian learning. *Journal of Neuroscience*, *16*, 1936–1947.

Mostofsky, S. H., Goldberg, M. C., & Landa, R. J. (2000). Evidence for a deficit in procedural learning in children and adolescents with autism: Implications for cerebellar contribution. *Journal of International Neuropsychological Society*, *6*, 752–759.

O'Reilly, R. C., & Frank, M. J. (2006). Making working memory work: A computational model of learning in the frontal cortex and basal ganglia. *Neural Computation*, *18*, 283–328.

O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. Cambridge, Massachusetts: MIT Press.

Ozonoff, S., & Jensen, J. (1999). Specific executive function profiles in three neurodevelopmental disorders. *Journal of Autism and Developmental Disorders*, *29*(2), 171–177.

Pascual-Leone, A., Wassermann, E. M., Grafman, J., & Hallett, M. (2004). The role of the dorsolateral prefrontal cortex in implicit procedural learning. *Experimental Brain Research*, *107*(3), 479–485.

Rinehart, N. J., Bradshaw, J. L., Brereton, A. V., & Tonge, B. J. (2001). Movement preparation in high-functioning autism and asperger disorder: A serial choice reaction time task involving motor reprogramming. *Journal of Autism and Developmental Disorders*, *31*, 79–88.

Rougier, N. P., Noelle, D. C., Braver, T. S., Cohen, J. D., & O'Reilly, R. C. (2005). Prefrontal cortex and flexible cognitive control: Rules without symbols. *Proceedings of the National Academy of Sciences*, *102*(20), 7738–7343.

Shah, A., & Frith, U. (1983). An islet of ability in autistic children: a research note. *Journal of Child Psychology and Psychiatry*, *24*(4), 613–620.

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, *28*, 643–662.

Tuchman, R., & Rapin, I. (2002). Epilepsy in autism. *Lancet Neurology*, *1*, 352–358.