A Subsymbolic Model of Language Pathology in Schizophrenia

Uli Grasemann **Risto Miikkulainen**
Department of Computer Sciences, The University of Texas at Austin
{uli, risto}@cs.utexas.edu

Ralph Hoffman
Department of Psychiatry, Yale University School of Medicine
ralph.hoffman@yale.edu

Abstract

This paper reports first results of a simulation of language pathology in schizophrenia. Using DISCERN, a subsymbolic model of story understanding and recall, the impact of different simulated lesions hypothesized to underlie schizophrenia is investigated. In response to excessive connection pruning, the model reproduces symptoms of delusions and disorganized language seen in schizophrenia, as well as the reduced output seen in compensated later states of the disorder. The effects of other lesions are less consistent with the symptoms of schizophrenia. The model therefore forms a promising basis for future computational investigations into the underlying causes of schizophrenia.

Keywords: Cognitive Modeling; Schizophrenia; Neural Networks; Natural Language Processing.

Introduction

Schizophrenia is a disabling psychiatric disorder characterized by complex alterations of language and thought. Establishing the physiological basis of its symptoms would greatly enhance our understanding of this disorder, possibly leading to more effective treatments. Yet after decades of intensive research, underlying brain processes remain uncertain (Kapur, 2003; Pantelis et al., 2005). Competing theories include reduced connectivity within and between cortical networks, drop-out of certain types of neurons, and disturbances involving subcortical dopamine (DA) neurons that affect cortical systems. Unfortunately currently available neuroimaging, pharmacological, and postmortem methods remain limited in characterizing neural substrates of schizophrenia. Thus, the schizophrenic brain remains a vexing "black-box" where output deviations challenge researchers to infer underlying processes. It is not surprising therefore that schizophrenia has become an important focus in the emerging field of computational models of psychopathology.

During the last two decades, researchers have tried to capture aspects of schizophrenia and many other disorders using connectionist models, with the goal of informing clinical research and resolving competing hypotheses about underlying brain mechanisms. In one of the earliest models of psychopathology, Hoffman (1987) investigated the emergence of "parasitic" stable states in attractor networks (Hopfield, 1982) to model delusions and hallucinations seen in schizophrenia. More recently, Ruppin, Reggia, and Horn (1996) used attractor networks in a computational model of Stevens's (1992) theory of the pathogenesis of schizophrenia. Spitzer (1997) and Silberman, Bentin, and Miikkulainen (2007) used self-organizing feature maps to model impaired lexical access in schizophrenia based on dysfunctional inhibition and priming mechanisms. Cohen and colleagues (Cohen & Servan-Schreiber, 1992; Braver, Barch, & Cohen, 1999) modeled the interactions between prefrontal cortex and the DA system to explain impaired processing of context in schizophrenia.

In a study more closely related to the present paper, Hoffman & McGlashan (1997, 2006) used simple recurrent networks (Harris & Elman, 1989) to simulate speech perception in order to understand the mechanisms underlying hallucinated speech, which is one characteristic symptom of schizophrenia. Excessive pruning of recurrent connections caused spontaneous speech percepts to be generated, thus emulating hallucinated speech. This result is particularly interesting because significant reductions in cortical connectivity occur ordinarily during adolescence. That schizophrenia usually first emerges during or shortly after this developmental period suggests that abnormal neurodevelopmental pruning may indeed be involved in the genesis of this disorder.

The study reported in this paper extends this hypothesis to three other critical manifestations of schizophrenia, specifically those that disrupt understanding and recall of stories. First, patients with delusions often discover spurious plots around them due to a tendency to confuse the actors in their personal stories with those of the shared stories of their culture. Secondly, patients with language disorganization tend to insert extraneous language material that derails story recall (Hoffman RE, Watts A, Varanko M, Lane D, Quinlan D, unpublished data). Finally, after an initial active stage of the disorder, patients tend to compensate for these positive symptoms, showing reduced symptoms at the cost of curtailed language output (Ventura et al., 2004).

The goal of this paper is to demonstrate that pathologies hypothesized to underlie schizophrenia can produce the abnormalities described above in a computational model of human story processing. The model used is an expanded version of DISCERN (Miikkulainen, 1993), a system that simulates human story understanding and recall using a multimodular neural network. DISCERN has been previously extended to include a memory architecture that permits the system to store and recall complex stories composed of multiple story segments, or scripts (Fidelman, Miikkulainen, & Hoffman, 2005). Further extensions of of the model include the capacity to process emotional content and an output filtering mechanism that allows the system to discard distorted language.

Simulations of several pathologies hypothesized to underlie schizophrenia were applied to DISCERN, including excessive connection pruning, distortion of episodic memory traces, and noise contamination of the system's working memory. The effects of these simulated lesions were then assessed to evaluate signs of (1) delusions, evidenced by a tendency to substitute one actor for another, (2) derailment-type disorganized speech, and (3) recovery from the active psychotic stage of the disorder at the cost of curtailed language output.

The main difference between the current work and previous approaches lies in the complexity of the model and the resulting richness of its observable behavior: using a connectionist model of story processing makes it possible to observe the behavioral effects of simulated lesions at the level of conversational discourse, creating a chance to further "bridge the gap between mind and brain, between clinical phenomena and underlying brain pathology" (Spitzer, 1997).

Script Theory

Script theory (Schank & Abelson, 1977) is one of the fundamental ideas underlying the model of human story processing used in the present study. It models the way humans process stereotypical sequences of events. For example, every time we walk into a restaurant, approximately the same thing happens: we wait to be seated, we order a meal and eat it, pay, then leave. The specific restaurant and the people with us may not be the same; certainly the price and quality of the food may change. The basic sequence of events, however, rarely does, and can therefore be learned and reused as a *script*.

Scripts are best understood as templates for certain types of situations, including open slots to be filled in (such as the kind of food), and constraints on what kinds of things can fill the slots (you cannot order the decor). An instance of a script, then, is a script whose slots have been filled to match a specific situation. In order to store and recall a specific event, all we need to do is remember the kind of script it followed and the concepts filling the slots.

Humans use scripts to interact efficiently with each other, grasp complex situations, and form expectations about a situation when faced with incomplete information. Scripts are central to the current understanding of human cognition, language, and memory. The hypothesis that humans use scripts is well supported by experimental evidence. For example, the degree to which events in a story will be remembered can be predicted by whether those events are part of a script (Graesser, Woll, Kowalski, & Smith, 1980). Similarly, the amount of time it takes humans to understand a sentence can be predicted by whether it fits into a script (Den Uyl & van Oostendorp, 1980). Script theory therefore forms a promising framework for computational models of story processing.

The Extended DISCERN Model

The original DISCERN system (Miikkulainen, 1993) was the first integrated subsymbolic model of human story processing. It consisted of several modules, each of which was responsible for one subtask of story processing. The model could understand and recall script-based stories and answer questions about them. The model was shown to exhibit important characteristics of human story processing, such as robustness to input errors and noise. However, it could only process stories that consisted of single script instances. Since realistic human stories usually consist of concatenations of several scripts, the model needed to be extended to multi-script stories.

This section provides an overview of the current extended version of DISCERN. In contrast to the original model, it is able to process multi-script stories of arbitrary length, process emotional context of stories, and filter overly distorted language output. A detailed account of an earlier version of the extended DISCERN was given by Fidelman et al. (2005). For simplicity, the extended model will be referred to as "DISCERN" in the rest of the paper.

Architecture Overview

Apart from the use of scripts as the basic unit of story understanding, the central concept of DISCERN is modularity. The task of understanding, recalling, and then paraphrasing a story is achieved by a chain of modules, each building on the results of the last module in the chain, and providing input for the next. The modules consist of simple recurrent or feedforward neural networks that are trained separately with backpropagation, then linked together to form the final system, as shown in Figure 1.

DISCERN reads and produces natural language. Each story consists of a sequence of scripts, but is presented to the system as plain text, one word at a time. While DISCERN understands and recalls the story, it is at different times represented at the level of words, sentences, scripts, and episodic memory traces. Figure 2 shows an example story and the representations used by DISCERN to encode the individual scripts. Each story in DISCERN is associated with an emotional context, represented as a pattern of neuron activations that encodes either positive, negative or neutral emotional tone. The emotion of a story plays an important role in story memory and recall, affecting the system's choice between alternative continuations of a story.

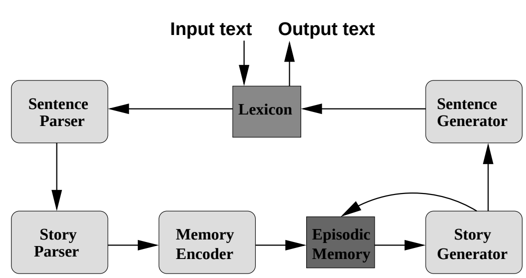

Figure 1: DISCERN is a neural network model of human story understanding and recall. The task of understanding and reproducing a story is achieved by a chain of modules, each building on the results of the previous module and providing input for the next.

The modules in DISCERN communicate using distributed representations of word meanings, i.e. fixed-size patterns of neuron activations, stored in a central lexicon. These representations are learned based on how the words are used in the example stories, using the FGREP algorithm (Miikkulainen, 1993), a modified version of backpropagation that treats input representations as an additional layer of weights.

A plain-text input story is first translated into input activations by the lexicon, then presented to the sentence parser one word at a time. The sentence parser builds a static representation of each sentence as it comes in. At the end of each sentence, that representation, which consists of a concatenation of word representations, is passed on to the story parser. The story parser in turn transforms a sequence of sentences into a static representation of a script, simultaneously building a representation of the story's emotional context. Script representations are called *slot-filler representations*, because they consist of a representation for the name for the script (the words starting with $ in Figure 2) and a sequence of concepts filling its slots.

At this point, the internal representation of a story consists of its emotional context and a list of slot-filler representations, one for each script in the story. The memory encoder turns this representation into episodic memory traces that can be successively recalled and reproduced by the story generator. This behavior is achieved using Recursive Auto-Associative Memory, or RAAM (Pollack, 1990), a neural network architecture that forms fixed-size distributed representations of recursive data structures like lists or trees.

Figure 3 shows the structure of the memory encoder. The network is trained to reproduce its own input, forcing it to form a compressed distributed representation of the input in its hidden layer. These representations are later used by the story generator as memory cues to access the episodic memory trace. Additionally, they are used as an encoding of the

```
Emotion: Negative
[$job Vito Mafia head likes New-York famous gangster]
Vito is a gangster .                   [Vito is _ _ gangster]
Vito is the head of the Mafia .        [Vito is Mafia _ head]
Vito works in New-York .               [Vito works New-York _ _]
Vito likes his job .                   [Vito likes _ his job]
Vito is a famous gangster .            [Vito is _ famous gangster]

[$driving Vito _ scared airport LA recklessly _]
Vito wants to go to LA .               [Vito wants LA goes _]
Vito enters his car .                  [Vito enters _ his car]
Vito drives to the airport .           [Vito drives airport _ _]
Vito is scared .                       [Vito is _ _ scared]
Vito drives recklessly .               [Vito drives _ _ recklessly]

[$pulled-over Vito cop arrests _ murder _ _]
Vito is pulled-over by a cop .         [Vito is cop _ pulled-over]
The cop asks Vito for his license .    [Cop asks license his Vito]
Vito gives his license to The cop .    [Vito gives cop his license]
The cop checks the license .           [Cop checks _ _ license]
The cop arrests Vito for murder .      [Cop arrests murder _ Vito]

[$trial Vito _ walks clears free murder good]
Vito is accused of murder .            [Vito is murder _ accused]
Vito is brought before the court .     [Vito is court _ brought]
Vito has a good lawyer .               [Vito has _ good lawyer]
The court clears Vito of murder .      [Court clears murder _ Vito]
Vito walks free .                      [Vito walks _ free _]
```

Figure 2: An example input story about a gangster getting arrested for a crime committed in another story. Each sentence is paired with the static representation used by DISCERN to represent it. Slot-filler representations of each script are also shown.

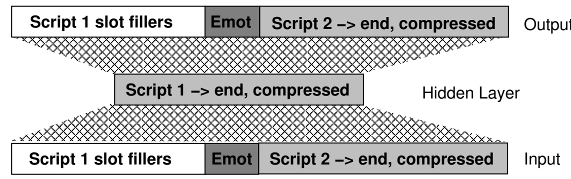

Figure 3: The memory encoder is a Recursive Auto-Associative Memory (RAAM) neural network (Pollack, 1990) that transforms a parsed story into episodic memory traces that can later be recalled by the story generator.

rest of the story by the memory encoder itself, while it steps backwards through each story, populating the episodic memory with script/memory-cue pairs.

With the memory traces in place, the system is now ready to recall the stories that were presented to it earlier. The story generator module, shown in Figure 4, is cued with the first memory in a story, then called repeatedly, producing a representation for a sentence each time, until it outputs a special "end of story" pattern. In addition to the next sentence, every cycle of the story generator produces a cue to the episodic memory that determines the next input. A memory cue consists of the compressed version of the rest of the current story, and can be thought of as a representation of the system's "discourse plan".

As long as the story generator produces sentences belonging to the same script, the memory cue does not change. Then, at the same time the last sentence of the script is produced, the cue changes, and the input is replaced by a memory of the next script. In this way, the story generator steps through each sentence of a story, and accesses each memory trace encoding it.

A recent extension to the original DISCERN is that the model may discard overly distorted language output. This is done by applying a filter every time information is retrieved from memory (either lexical or episodic). For example, if a memory cue output by the story generator is sufficiently different from all cues in episodic memory (measured by Euclidean distance), DISCERN discards the memory, instead ending the story. Similarly, if a word representation produced by the system does not have a close enough match in the lexicon, the word is discarded.

Finally the sentence generator, last in the chain, takes the sentence representations produced by the story generator and turns them back into a sequence of individual words. The system then outputs plain text translations of these words as provided by the lexicon.

For the present study, only the behavior of the memory encoder and the story generator are of immediate interest. Furthermore, when examining the effects of lesions on the story generator, it is desirable to separate the effects of errors that are due to the lesioning from those that are not. Therefore, the parsing modules and the sentence generator were omitted from the simulations, focusing the analysis on the lexicon, the memory encoder, and the story generator.

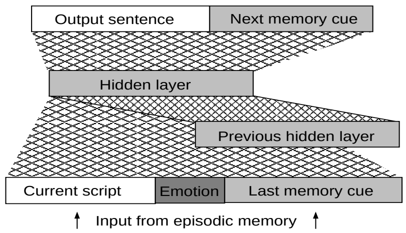

Figure 4: The story generator in DISCERN is a simple recurrent neural network that reproduces a story by successively recalling each episodic memory trace encoding it. While recalling a story, it produces both the next sentence and a memory cue that determines its own next input. In this manner, it can reproduce stories consisting of an arbitrary number of scripts.

Experiments

In the experiments reported below, different levels of simulated lesions hypothesized to underlie schizophrenia were applied to DISCERN, with the goal of reproducing the documented symptoms of the disorder. This section describes the input data used, the lesions applied to the model, and the experimental methods and results.

Input data

In order to give the system an opportunity to fail in interesting ways, a corpus of 28 input stories was developed that was much more extensive and complex than others previously used with DISCERN. Stories ranged between three and seven scripts long and were divided roughly into two groups: the first group described normal occurrences, mostly concerning the life of a "Self" character – a person that was overrepresented in the corpus to simulate the concept of the person telling the stories. This part of the corpus included stories with a negative emotional tone, such as the "Self" character driving drunk and getting caught by the police, as well as stories about positive events, like visiting relatives, or the self character being praised by his boss. The second group of stories consisted of mostly negative stories about a group of gangsters going about their gangster business – committing crimes, killing each other, and occasionally getting caught. Figure 2 shows an example.

All stories were assembled from 14 different scripts describing stereotypical sequences of events such as meeting someone for a drink or being pulled over by the police. Overall, the corpus contained 550 single sentences in 120 script instances. The lexicon contained 170 words, including 20 names or descriptions of characters in the stories (e.g. "Frank" or "lawyer").

Methods

The first step in the experiments was to develop the word representations used in the lexicon. Since sentence processing networks are generally best suited for the task of producing word representations, the sentence parser and the sentence generator were trained using FGREP for 5000 iterations of the entire corpus. Each word representation consisted of a fixed-size pattern of 12 neuron activations. The network learning rate was initially set to 0.1, then reduced exponentially with a half-life of 3000 iterations. The training rate for the word representations was always 5% of the network training rate.

By the end of training, the word representations had converged to good semantic representations of the concepts. Words whose representations were similar tended to denote similar concepts, and usually belonged to the same word category. Names of story characters, for instance, formed a tight and well-defined cluster: with only a single exception, the five words closest to each of the ten names were either another name, or the word "man". The trained sentence processing networks themselves achieved close to perfect performance on the sentences in the corpus.

With the word representations in place, 10 DISCERN systems consisting of a memory encoder and a story generator were then trained for 10,000 iterations. The initial training rate was again set to 0.1, and was reduced with a half-life of 3000 epochs. The training algorithm was standard backpropagation, so the word representations were not changed any further. The hidden layer of each memory encoder consisted of 48 neurons, and the story generators had 150 hidden neurons. All parameters were set empirically.

Using the lexicon and the trained networks, the impact of several simulated lesions hypothesized to underlie schizophrenia was then evaluated in detail. The lesions included connection pruning, noise contamination of working memory, and distortion of episodic memory traces.

Connection pruning was modeled in DISCERN by removing any connections between the recurrent and hidden layers of the story generator whose absolute weights were below a threshold. By increasing the threshold, increasing percentages of connections were cut.

Noise contamination of working memory, intended to model impaired processing of context possibly due to dopamine imbalance, was done by adding increasing levels of noise to the recursive layer of the story generator during each cycle. Similarly, episodic memory traces were distorted by adding increasing levels of noise to the cues produced by the memory encoder network.

As the level of each lesion increased, the performance of the system was recorded, including the amount of language output, the number of stories in which derailments occurred, and the number and type of word substitutions. All reported results are averages over all 10 systems.

Results

The effects of connection pruning were found to model the behavioral changes seen in schizophrenia more closely than the other lesions investigated. The results reported in this section therefore focus on pruning as the central pathology. Figure 5 summarizes the behavioral changes as a function of the number of connections cut. Figure 5a shows the percentage of derailed stories, i.e. stories in which the system retrieves a faulty episodic memory and jumps to another story. Figure 5b shows the number of times the system substituted one word for another. It also shows the number of errors within the same word class (e.g. noun → noun), and the number of times one actor was substituted for another (e.g. "Frank" → "Joe"). Finally, Figure 5c plots the reduction in language output due to the output filtering mechanism described earlier.

In the initial unpruned state (left edge of the plots), the story recall performance of all systems is close to perfect: DISCERN never derails from one story into another, and only rarely produces word substitutions. The output filtering leads to no reduction in language production.

At moderate levels of pruning, both derailments and actor substitutions increase sharply as a function of pruning strength, suggesting positive psychotic symptoms consistent with the initial active phase of schizophrenia. At the peak of positive symptoms, language production is only moderately reduced by the output filter.

At high levels of pruning, filtering reduces both derailments and actor substitutions dramatically, at the cost of highly curtailed language output. This is consistent with a possible compensatory function of negative symptoms common in later stages of schizophrenia, suggesting that in compensated states of the disorder, the underlying disorganization of language processing remains but is rarely expressed.

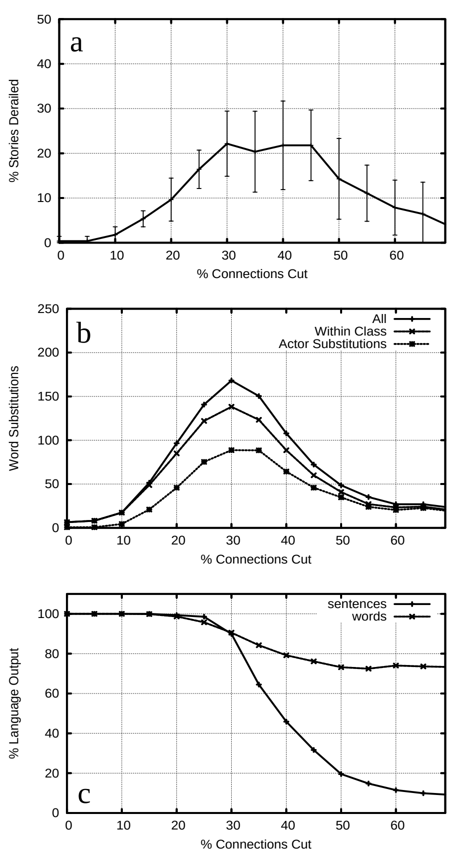

Figure 5: The impact of a simulated lesion (connection pruning) on (a) the percentage of derailed stories, (b) the number and type of word substitutions, and (c) the amount of language produced. Both derailment and actor substitutions increase sharply in response to pruning, consistent with the positive symptoms during the initial psychotic phase of schizophrenia. At higher levels of pruning, output filtering of the story generator reduces these symptoms at the cost of curtailed language output, modeling a compensatory role of negative symptoms as the condition progresses.

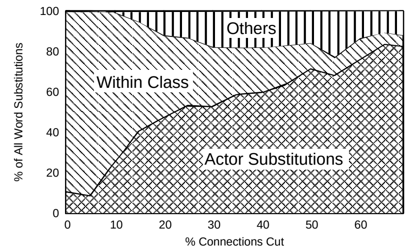

Figure 6: The types of word substitutions produced in response to connection pruning. In the unlesioned system, only a small percentage of word substitutions exchange one actor for another. In the pruned system, most word substitutions involve actors, supporting the interpretation of actor substitutions as signs of delusions.

Figure 6 illustrates the relative frequencies of different types of word substitutions observed. In the unpruned system, word substitutions are not only rare – they also do not tend to involve actors, which supports the interpretation of actor substitutions as signs of delusions. As the level of pruning increases, actor substitutions are responsible for an increasing percentage of overall errors. Furthermore, substitutions rarely cross word categories (e.g. noun → verb), which is consistent with clinical observations.

A further interesting finding concerns the structure of the actor substitutions. A delusional patient would be expected to not just randomly exchange one actor for another, but would instead tend to consistently insert specific actors from other stories. Furthermore, patients with delusions tend to insert themselves into shared stories of their culture. The actor substitutions observed in the pruned system show both of these characteristics: at the peak level of positive symptoms, a group of three actors accounts for 50–70% of the actor substitutions in all ten systems. In six of the ten systems, the "Self" character is part of that group.

The effects of both noise contamination of working memory and distortion of memory traces (not shown) are less consistent with the symptoms seen in schizophrenia. For example, both lesions produce unrealistically high levels of derailment ($\approx$ 50%), and do not tend to insert the "Self" character.

Discussion and Future Work

The results demonstrate that DISCERN provides a good model of delusions and the type of language disorganization exhibited by patients with schizophrenia. In addition to these positive psychotic symptoms suggestive of the active phase of schizophrenia, the model develops curtailed language output at increased levels of lesioning, modeling the negative symptoms that commonly follow episodic positive symptom exacerbations.

In patients, these negative symptoms appear to be more enduring and progress over the lifetime of the individual (Ventura et al., 2004). The DISCERN model provides a ready explanation for this pattern – output filtering dramatically curtails errors, at the cost of reducing overall language output.

Our results suggests that underlying disorganization remains in individuals in the compensated phase – it is just not expressed.

Future work will focus on using the model to gain insight into the physiological causes of schizophrenia. Further extension of the model will make it possible to include additional candidate pathologies and refine the models of the lesions already investigated. Neurodevelopmental pruning, for example, could be more closely modeled by eliminating connections continuously throughout the training process. More sophisticated models of output filtering based on altered neural response characteristics may provide a better characterization of compensatory reactions as well as model effects of antipsychotic medication. Further possible lesions include cell death, intermodular disconnection, noise contamination of semantic representations, and altering response characteristics of neurons to mimic neuromodulator effects of dopamine.

Another intriguing possibility is to use the model to simulate disorders other than schizophrenia. For example, lesions that tend to cause the system to jump from one coherent story to another may be indentified, thus providing a model for manic speech behavior. Such behavior is distinct from the language disorganization in schizophrenia, where fragments from multiple stories are thrown together. In this manner the model can be used not only to investigate possible causes of schizophrenia, but also to distinguish it from other disorders.

Conclusions

This paper reported first results of a simulation of language pathology in schizophrenia. Using an extended version of the DISCERN neural network model of story understanding and recall, the impact of different levels of simulated pathologies hypothesized to underlie schizophrenia was investigated. In response to excessive connection pruning, the model reproduced symptoms of delusions and disorganized language seen in schizophrenia, as well as the reduced output seen in compensated later stages of the disorder. The effects of other lesions were less consistent with the symptoms of schizophrenia. The model therefore forms a promising basis for future computational investigations into the underlying causes of schizophrenia.

Acknowledgments

This work was supported by NIMH under grant R01MH066228 and by NSF under grant IIS–0083776. The authors would like to thank Peggy Fidelman and Hong Ming Yeh for their previous work on extending DISCERN.

References

Braver, T., Barch, D., & Cohen, J. (1999). Cognition and control in schizophrenia: A computational model of dopamine and prefrontal function. *Biol. Psychiat.*, *46*(3), 312-28.

Cohen, J., & Servan-Schreiber, D. (1992). Context, cortex and dopamine: A connectionist approach to behaviour and biology in schizophrenia. *Psychol. Rev.*, *99*, 45–77.

Den Uyl, M., & van Oostendorp, H. (1980). The use of scripts in text comprehension. *Poetics*, *9*, 275–94.

Fidelman, P., Miikkulainen, R., & Hoffman, R. (2005). A subsymbolic model of complex story understanding. In *Proc. of CogSci 2005*.

Graesser, A., Woll, S., Kowalski, D., & Smith, D. (1980). Memory for typical and atypical actions in scripted activities. *J. Exp. Psychol.-Hum. L.*, *6*, 503–15.

Harris, C., & Elman, J. (1989). Representing variable information with simple recurrent networks. In *Proc. of CogSci 1989* (pp. 635–42). Hillsdale, NJ: Erlbaum.

Hoffman, R. (1987). Computer simulations of neural information processing and the schizophrenia-mania dichotomy. *Arch. Gen. Psychiat.*, *44*, 178-88.

Hoffman, R., & McGlashan, T. (1997). Synaptic elimination, neurodevelopment and the mechanism of hallucinated 'voices' in schizophrenia. *Am. J. Psychiat.*, *154*, 1683–9.

Hoffman, R., & McGlashan, T. (2006). Using a speech perception neural network computer simulation to contrast neuroanatomic versus neuromodulatory models of auditory hallucinations. *Pharmacopsychiat.*, *39 (suppl. 1)*, 554–64.

Hopfield, J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci.*, *79*, 2554–58.

Kapur, S. (2003). Psychosis as a state of aberrant salience: a framework linking biology, phenomenology, and pharmacology in schizophrenia. *Am. J. Psychiat.*, *160*(1), 13–23.

Miikkulainen, R. (1993). *Subsymbolic natural language processing: An integrated model of scripts, lexicon, and memory*. Cambridge, MA: MIT Press.

Pantelis, C., Yucel, M., Wood, S., Velakoulis, D., Sun, D., Berger, G., et al. (2005). Structural brain imaging evidence for multiple pathological processes at different stages of brain development in schizophrenia. *Schizophrenia Bull.*, *31*(3), 672–96.

Pollack, J. (1990). Recursive distributed representations. *Artificial Intelligence*, *46*(1), 159–216.

Ruppin, E., Reggia, J., & Horn, D. (1996). Pathogenesis of schizophrenic delusions and hallucinations: A neural model. *Schizophrenia Bull.*, *22*(1), 105–23.

Schank, R., & Abelson, R. (1977). *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures*. Hillsdale, NJ: Erlbaum.

Silberman, Y., Bentin, S., & Miikkulainen, R. (2007). Semantic boost on episodic associations: An empirically based computational model. *Cognitive Sci.* (In Press)

Spitzer, M. (1997). A cognitive neuroscience view of schizophrenic thought disorder. *Schizophrenia Bull.*, *23*(1), 29–50.

Stevens, J. (1992). Abnormal reinnervation as a basis for schizophrenia: A hypothesis. *Arch. Gen. Psychiat.*, *49*, 238–43.

Ventura, J., Nuechterlein, K., Green, M., Horan, W., Subotnik, K., & Mintz, J. (2004). The timing of negative symptom exacerbations in relationship to positive symptom exacerbations in the early course of schizophrenia. *Schizophrenia Res.*, *69*(2–3), 333–42.