

Hyperlearning: A Connectionist Model of Psychosis in Schizophrenia

Uli Grasemann

Department of Computer Sciences, The University of Texas at Austin
{uli, risto}@cs.utexas.edu

Risto Miikkulainen

Ralph Hoffman

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

Abstract

Abnormal brain processes that underlie schizophrenia are incompletely understood. Diagnosis of this disorder relies in large part on psychotic symptoms that are observed through conversational language. In this paper, two such symptoms, delusions (fixed false beliefs) and derailments (inability to follow a coherent discourse plan) are modeled using DISCERN, a connectionist model of human story processing. Simulations of alternative pathologies thought to underlie schizophrenia are applied to DISCERN, and the resulting language abnormalities are evaluated for symptoms of schizophrenia. “Hyperlearning”, a simulation of excessive dopamine release, is shown to produce a compelling model for both delusional and derailed language. Applied to different locations in the model, hyperlearning led to different symptoms, suggesting how clinical subtypes of schizophrenia could arise from a common underlying process.

Keywords: Neural Networks; Schizophrenia; Natural Language Processing.

Introduction

No current laboratory or imaging technique can reliably identify individuals with schizophrenia. Instead, diagnosis of this disorder relies on symptoms observed in clinical interviews using conversational language. The symptoms of schizophrenia are complex and span a wide range of altered behavior and perception. The present study focuses on delusions and derailment, two psychotic symptoms that play an important role in diagnosing patients with schizophrenia.

Delusions, a major characteristic of the “paranoid type” of schizophrenia, can take many different forms. In general, delusions are fixed false beliefs that cannot be changed through rational refutation. Delusions in schizophrenia occur in more than 60% of patients, and often have grandiose or paranoid content (Harrow et al., 2004; Applebaum, Clark Robbins, & Roth, 1999), i.e. a patient might believe that he is the lost son of John Lennon, or that he is being monitored by the CIA. In this way, delusions in schizophrenia often insert the self or persons close to the patient into rigid, implausible or bizarre narrative schemata.

Derailment behavior, the second focus of the present study, is fluent speech that fails to adhere to an organizing topic or frame of reference (Bleuler, 1911; Andreasen, 1979). When telling a story, for example, a patient may repeatedly skip to other unrelated stories, leaving the listener in a bewildered state. Derailments often occur in “disorganized type” schizophrenia, where symptoms are dominated by disorganized language and behavior.

Establishing the brain processes that lead to delusions and derailment in schizophrenia would greatly advance our understanding of this disorder. Yet after almost a century of

clinical research, the brain processes underlying these and other symptoms of schizophrenia remain incompletely understood (Harrison, 1999). Only recently, theories have begun to emerge that have the potential to explain how abnormalities at the brain level might lead to the emergence of symptoms.

One recent theory, put forth by Kapur (2003), is based on the view that dopamine (DA) mediates the significance, or “saliency”, of subjective experience (Berridge & Robinson, 1998). Kapur proposes that in schizophrenia, an overabundance of midbrain DA leads to a pathological enhancement of saliency, which in turn causes psychotic symptoms. Delusions, for example, are explained as secondary reactions to an altered experience of the world — i.e. as an attempt by the brain to make sense of the excessive significance assigned to insignificant events. The theory that some symptoms of schizophrenia are caused by excessive saliency has been widely endorsed in the psychiatric literature, and forms the theoretical basis of the present computational study.

Excessive saliency (and thus excessive DA release) was simulated in a neural network model of human story processing using artificially high network learning rates. In contrast to other network disturbances, this “hyperlearning” simulation was shown to produce a compelling model for both delusional and derailed language.

Approach

The approach taken in the present study is based on DISCERN (Miikkulainen, 1993; Fidelman, Miikkulainen, & Hoffman, 2005; Grasemann, Miikkulainen, & Hoffman, 2007), a neural network model of human story processing. Simulations of alternative pathologies thought to underlie schizophrenia, including “hyperlearning”, a model of excessive saliency, were applied to the model. The resulting language disturbances in DISCERN were then evaluated for characteristic signs of delusions and derailments in schizophrenia.

Several computational studies in the past were based on similar approaches. Hoffman (1987) and Ruppín, Reggia, and Horn (1996), for example, propose models of psychotic symptoms based on attractor networks. Cohen and colleagues (Cohen & Servan-Schreiber, 1992; Braver, Barch, & Cohen, 1999) focus on modeling behavioral deficits and cognitive impairment based on a neural network-based model of the frontal cortex. Spitzer (1997) uses self-organizing maps to simulate aspects of disturbed lexical access in schizophrenia. Finally, Hoffman and colleagues (1997; 2006) investigate mechanisms by which hallucinations can arise in speech perception networks.




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

The most important way in which DISCERN differs from these previous models is the level of behavior at which symptoms are observed. Using a model of human story processing makes it possible to observe the effects of network disturbances at the level of conversational language — the same level at which the clinical symptoms that define schizophrenia are diagnosed.

The remainder of this section first gives a brief overview of DISCERN, and then discusses the network disturbances used to model possible causes of schizophrenic symptoms.

The DISCERN Model

DISCERN is a neural network-based model of human story understanding and recall. Stories in DISCERN are sequences of *scripts* – stereotypical event sequences that can be adapted to match specific situations. A classic example for a script is visiting a restaurant: The sequence of events (i.e. waiting to be seated, ordering, etc.) usually stays the same, and can be adapted to match a specific restaurant visit by filling in open *slots* such as the type of restaurant or the kind of food. Using scripts, specific events (called script *instances*) can be understood and remembered by storing only the type of script they follow and the concepts filling the script’s slots.

In DISCERN, 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 with backpropagation and 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 in 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 memories. 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.

```
[ $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 earlier. The story consists of four scripts: the slots of each script are on top, followed by the sentences of the script (left) and their static representations used by DISCERN (right). During story understanding, DISCERN translates such stories from individual words to sentence representations, slot-based script representations, and finally episodic memory traces. Story recall reverses this process.

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 (Mikkulainen, 1993), a modified version of backpropagation that treats input representations as an additional layer of weights.

Words in the input text are first presented to the lexicon, which translates them into word representations. The representations are then passed on 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 a string of episodic memories 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.

The representations of script sequences produced by the memory encoder are later used by the story generator as memory cues to address the episodic memory.

With the episodic memory in place, the system is now ready to recall the stories that were presented to it earlier. The story generator module is cued with the first memory in each 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.

While the story generator produces sentences belonging to the same script, the memory cue stays the same. Then, at the same time the last sentence of a 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 recalls each memory encoding it.

Finally the sentence generator, last in the chain, takes the static 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.

Modeling Impaired Story Processing

DISCERN’s architecture is complex and provides a wide range of opportunities to simulate pathological brain processes and observe their effects. The main focus of this work is Kapur’s (2003) theory that psychotic symptoms in schizophrenia are the result of excessive salience.

Within Kapur’s framework, excessive salience is assumed to be caused by increased midbrain DA release. However, since the theory does not build on a detailed physiological model, brain processes by which increased salience may in turn cause psychotic symptoms remain necessarily abstract: Symptoms are simply the result of the brain trying to adapt to an altered experience of the world.

The present paper extends Kapur’s theory by adding “hyperlearning”, a specific mechanism that could link excessive DA release and psychotic symptoms directly, and may serve as a possible neural substrate for excessive salience. Hyperlearning simulates excessive DA release in DISCERN by performing additional backpropagation learning with artificially increased learning rates.

The approach is motivated by two pieces of evidence linking excessive DA release, psychotic symptoms, and overly intense learning in humans. First, excessive midbrain DA release is likely to occur during some stages of schizophrenia (Laruelle, 2000), and has been linked to increased learning intensity (Gibbs et al., 2007). Second, a tendency to mismatch expectancy of observed experience and outcome has been associated with delusions in patients with schizophrenia (Corlett et al., 2007); elevated prediction errors of this kind have been shown to produce a “super-learning” state in human subjects (Aitken, Larkin, & Dickinson, 2000).

Intuitively, salience reflects a chance to gain meaningful knowledge from an experience. Given that increased learning intensity is linked to both excessive DA and psychotic symptoms, it is a reasonable assumption that hyperlearning

in DISCERN may capture important aspects of the impact of excessive salience. In the experiments reported below, hyperlearning was therefore used as a first approximation to modeling Kapur’s theory.

Hyperlearning was applied separately to DISCERN’s memory encoder network and to its story and sentence generators, reflecting the possible impact of increased salience on episodic memory encoding and recall.

Additionally, a range of other network disturbances were investigated for comparison. First, loss of cortical connectivity was modeled in DISCERN by removing connections in the story generator or memory encoder networks whose absolute weights were below a threshold. Following the theory that symptoms in schizophrenia may be the result of brain processes that try to compensate for lost connectivity, pruning was investigated both by itself and in combination with additional backpropagation learning.

Noise contamination of working memory, intended to model impaired processing of context associated with schizophrenia, was simulated by adding increasing levels of noise to the context layer of the story generator during each cycle.

Finally, impaired lexical access, thought to underlie language disorganization in schizophrenia, was modeled by adding noise to the word representations in the lexicon.

Experiments

The experiments reported below are based on a set of 28 hand-coded input stories. The stories ranged between three and seven scripts long and were divided into two groups: The first group described normal experiences in the daily life of a “Self” character, a person that was overrepresented in the corpus to simulate the one experiencing and reproducing 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 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 approximately 550 single sentences in 120 script instances. The lexicon contained about 170 words, including 20 names or descriptions of characters in the stories (e.g. “Joe” or “lawyer”).

The first step in the experiments was to develop word representations to be used in the lexicon. A sentence parser network was trained using the FGREP algorithm to obtain the word representations. Each word representation consisted of a pattern of 12 neuron activations. After 500 iterations of the entire corpus, the word representations had converged to good semantic representations of the concepts. Similar word representations tended to stand for similar concepts, and usually belonged to the same word category. Names of story char-

acters, 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”.

With the word representations in place, ten complete DISCERN systems were trained to reproduce all 28 input stories. Network sizes and training parameters were determined empirically. Sentence parsers and generators had 250 hidden neurons, story parsers had 225, and story generators had 150. Memory encoder networks had 48 hidden neurons.

All modules were first trained separately until they achieved a reasonable performance, and then linked in a chain and trained for an additional 10,000 epochs for fine-tuning. The training algorithm was standard backpropagation, so the word representations were not changed any further. The learning rate for each module was recalculated after each iteration of the story corpus by multiplying the module’s average output error by a constant factor (0.4). In this way, the learning rate was automatically decreased during the training process to allow for more accurate training.

After training, the resulting DISCERN systems were able to reproduce all 28 stories almost perfectly. The average recall performance (measured as the percentage of sentences correctly reproduced) ranged from 95% to 97%. Errors usually involved substituting one word for another that was used in a very similar way in the input stories. Two of the ten systems switched once from one story to another, but both times, the switch was to a closely related story. This type of “benign” derailment is not uncommon in healthy individuals (Hoffman, Stopek, & Andreasen, 1986).

Each of the lesions discussed previously was applied to all ten healthy DISCERN systems. The intensity of each lesion was increased until a realistic reduction of recall performance was reached (about 50% of healthy performance). Statistics on recall performance, number and type of lexical errors, and derailment behavior were collected for each lesion.

The results were then evaluated to determine whether or not the models produced language disturbances consistent with schizophrenia. Specifically, a realistic model would be expected to generate grammatically correct language with low but non-zero levels of lexical errors, while at the same time showing signs of delusions and/or derailments.

Derailments were scored when DISCERN switched from one story context to another. Plausible derailment behavior should contain frequent jumps between dissimilar stories.

Delusions are harder to quantify than derailments, but should involve recurring and well-formed new narratives that do not occur in the original story corpus. In patients, delusions are associated with “agency shifts”, i.e. migrations of characters from one story context to another. Often, delusional stories involve the self being inserted into implausible or bizarre narrative schemata.

Results

Of all lesions investigated, only hyperlearning lead to plausible simulations of delusional language. Applied to the story and sentence generator networks, hyperlearning robustly pro-

duced stable patterns of “agency shifts” where characters migrated between stories and produced meaningful new narratives. These agency shifts often involved the “self” character migrating to gangster stories and other violations of global context, while the local story structure stayed intact. The following text is an example from such a story.

```
[...]  
*I(Tony) was scared.  
*I(Tony) drove carefully.  
*I(Tony) entered City-Hall for bombing.  
*I(Tony) bomb(ed) City-Hall.  
The *wedding(bombing) was a success.  
*I(Tony) made a phone-call.  
*I(Tony) smoked a cigarette.  
[...]
```

Here, the “self” consistently replaces the gangster Tony. Apart from this agency shift, the word “wedding” intrudes from another context to replace the word “bombing. Patterns like these often recurred many times in the output of a single DISCERN system, creating a compelling simulation of delusional narratives.

Agency shifts in this simulation dominated other types of errors. Consistent with language behavior in patients with delusions, the syntax of the output language remained intact, and levels of lexical errors besides agency shifts were low but non-zero. Derailments following this type of hyperlearning were rare, which is consistent with the “delusional type” of schizophrenia, where symptoms do not tend to include prominent language disorganization.

Hyperlearning did, however, produce frequent derailments when applied to the memory encoder instead of the story and sentence generators. The following output was produced by a DISCERN system after 500 epochs of hyperlearning:

```
[...]  
Vince went to Starbucks.  
[Derailing to story 13]  
I went to Starbucks.  
*Vince(I) sat at a table.  
I ordered coffee.  
I drank the coffee.  
*Vince(I) met *Joe(Mary) at Starbucks.  
[Derailing to story 15]  
Mary was the girlfriend of Joe.  
[Derailing to story 12]  
*Vito(Joe) was the fiancée of *Tony(Mary).  
I hated Joe.  
*Vince(I) distrusted Joe.  
[...]
```

The language in this case seems clearly disorganized, and the system becomes unable to follow a coherent story line. This example is particularly interesting because derailments are accompanied by intrusions of characters from the original story context, suggesting that the disorganization occurs at a deeper level than just that of distorted memory retrieval. Other examples where word substitutions seemed to serve as a segway into derailed discourse appear to confirm this view:

```
[...]  
I was drunk.  
I drove recklessly.  
I was pulled-over by a cop.  
The cop asked *Vince(me) for *his(my) license.  
[Derailing to story 25]  
Vito was pulled-over by a cop.
```




Figure 3: The average number of unique vs. repeated character substitutions resulting from (A) hyperlearning of generator modules, (B) semantic noise, (C) connection pruning, and (D) working memory noise. As the intensities of the network disturbances increase, all produce more and more character substitutions. However, the substitutions following hyperlearning are by far the most consistent, leading to the emergence of stable patterns of “agency shifts” that resemble delusional narratives.

the cop asked Vito for his license.
 Vito gave his license to the cop.
 the cop checked the license.
 the cop arrested Vito for bombing.
 [...]

In this case, the switch to a gangster story is preceded by the intrusion of a gangster character. Note also how DISCERN adjusts the pronoun to match the new subject.

Not all instances of derailed language were equally consistent with disorganized language in schizophrenia. In particular, DISCERN tended to oscillate between two stories, which occurs only rarely in patients with schizophrenia. Future refinements of the model will address this issue.

To summarize the results so far: The hyperlearning simulation of aberrant salience led to plausible simulations of both delusions and derailed language. Different symptoms emerged depending on the subset of modules to which hyperlearning was applied, suggesting a possible mechanism by which different clinical subtypes of schizophrenia could emerge from a common brain process.

The remainder of this section compares hyperlearning to the other lesions applied to the model. First, as mentioned above, the simulation of delusional language was unique to hyperlearning – no other lesion led to patterns of agency shifts that were as stable, or tended to insert the “I” character as frequently. Figure 3 compares the consistency of agency shifts across different lesions, i.e. how many times specific substitutions of story characters were repeated within the same story. Substitutions following hyperlearning are much more likely to be repeated, leading to the consistent “delusional” patterns observed.

Derailments, on the other hand, did occur frequently following several lesions, including pruning and working memory noise. However, with lesions that did not include additional training, derailments were accompanied by high levels of ungrammatical language and lexical errors, both of which are inconsistent with the kind of language seen in schizophrenia (Figure 4). The following text, for example, was produced by a DISCERN system while noise was added to the story generator’s working memory:

[...]
 *Kate(Stacy) was from New-York.
 *Mary(Stacy) *smoked(drove) a *praised(compact).
 *Mary(Stacy) liked *baseball(movies).
 [Derailing to story 22]
 Vince talked to Vito about guns.
 *Tony(Vince) *kissed(liked) guns.
 Vince talked to *Bob(Vito) a long time.
 Vince liked talking to *Tony(Vito).
 Vince gave a *hand-shake(kiss) good-bye to *man(Vito).
 [...]

The system does derail to another story, but sentences like “Mary smoked a praised” and “Tony kissed guns” make the output random and non-sensical rather than disorganized.

Finally, when pruning was applied to the memory encoder during training, DISCERN did produce derailments without ungrammatical language or frequent lexical errors. However, over 90% of network connections had to be cut in order to approximate a realistic reduction in recall performance, which makes this kind of pruning an unlikely candidate pathology. Furthermore, since pruning does not lead to a model for delusions, two separate processes would have to be assumed unnecessarily for derailments and delusions.

Additional results, including the full output text of all ten DISCERN systems for all lesions investigated, can be found at <http://nn.cs.utexas.edu/?schizo>.

Discussion and Future Work

The results demonstrate that hyperlearning provides plausible simulations of both delusions and derailment behavior in schizophrenia. The present study therefore supports Kapur’s (2003) theory of increased salience, and furthermore offers a computational account of a brain process through which abnormal salience could lead to psychotic symptoms.

Additionally, the model suggests a possible refinement, namely that different clinical subtypes of schizophrenia may arise depending on the impact of overly intense learning on different functional modules. These findings demonstrate how computational models of psychopathology can complement clinical research by creating and formulating hypotheses about the link between mind and brain, between underlying brain processes and their behavioral manifestations.

Other possible pathologies underlying derailment and delusions in schizophrenia, including working memory impairment, loss of cortical connectivity, and disturbed lexical




Figure 4: Lesions without additional training lead to high levels of ungrammatical language as recall performance declines. In contrast, hyperlearning produces almost no ungrammatical language.

access, were not supported by simulations. However, DA imbalance in schizophrenia is likely a secondary effect of other abnormal processes, and future research may provide insight into the role of further pathologies involved in schizophrenia.

The present model focuses on delusions and derailments, and does not account for other symptoms of schizophrenia. Specifically, negative symptoms (as opposed to psychotic symptoms) are not addressed. Negative symptoms include reduced language output and flattened emotions. In future work, emotional intensity will in part determine the intensity of the learning process, and the possible role of emotional flattening as a compensatory mechanism will be investigated.

A further symptom that is not addressed at present are hallucinations, which primarily occur as spoken speech in schizophrenia. Within the framework of abnormal salience, hallucinations are assumed to arise from “the abnormal salience of the internal representations of percepts and memories.” (Kapur, 2003). However, a previous computational model suggests that connection loss in speech perception networks similar to those in DISCERN can model hallucinated speech (Hoffman & McGlashan, 1997). Future work will attempt a unified account of hallucinations in schizophrenia.

Conclusion

A computational model of language disturbance in schizophrenia was proposed and evaluated. Using DISCERN, a neural network-based model of human story processing, different brain abnormalities thought to underlie schizophrenia were simulated and compared with respect to their ability to model symptoms of schizophrenia. “Hyperlearning”, a simulation of Kapur’s (2003) theory of abnormal salience as the basis for psychosis in schizophrenia, was shown to provide a plausible model of delusions and derailment behavior, two key symptoms of schizophrenia.

Acknowledgments

This work was supported by NIMH under grant R01MH066228 and by NSF under grant IIS-0083776.

References

- Aitken, M. R., Larkin, M. J., & Dickinson, A. (2000). Super-learning of causal judgments. *Q J Exp Psych*, *53*, 59–81.
- Andreasen, N. C. (1979). Thought, language, and communication disorders. II. diagnostic significance. *Arch Gen Psychiat*, *26*, 1325–30.
- Applebaum, P. S., Clark Robbins, P., & Roth, L. H. (1999). Dimensional approach to delusions: comparison across types and diagnoses. *Am J Psychiat*, *156*, 1938–43.
- Berridge, K., & Robinson, T. (1998). What is the role of dopamine in reward: hedonic impact, reward learning, or incentive salience? *Brain Res Brain Res Rev*, *28*, 309–69.
- Bleuler, E. (1911). Dementia praecox oder die Gruppe der Schizophrenien. In G. Aschaffenburg (Ed.), *Handbuch der Psychiatrie*. Leipzig: Breitkopf und Hartel.
- 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.
- Corlett, P., Murray, G., Honey, G., Aitken, M., Shanks, D., Robbins, T., et al. (2007). Disrupted prediction-error signal in psychosis: evidence for an associative account of delusions. *Brain*, *130*, 2387–400.
- Fidelman, P., Miikkulainen, R., & Hoffman, R. E. (2005). A subsymbolic model of complex story understanding. In *Proc of CogSci’05*.
- Gibbs, A. A., Naudts, K. H., Spencer, E. P., & David, A. S. (2007). The role of dopamine in attentional and memory biases for emotional information. *Am J Psychiat*, *164*, 1603–9.
- Grasemann, U., Miikkulainen, R., & Hoffman, R. E. (2007). A subsymbolic model of language pathology in schizophrenia. In *Proc of CogSci’07*. Nashville, Tennessee, USA.
- Harrison, P. (1999). The neuropathology of schizophrenia a critical review of the data and their interpretation. *Brain*, *122*, 593–624.
- Harrow, M., Herbener, E. S., Shanklin, A., Jobe, T. H., Rattebury, F., & Kaplan, K. J. (2004). Followup of psychotic outpatients: dimensions of delusions and work functioning in schizophrenia. *Schizophrenia Bull*, *30*, 147–61.
- 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.
- Hoffman, R., Stopek, S., & Andreasen, N. (1986). A comparative study of manic versus schizophrenic speech disorganization. *Arch Gen Psychiat*, *43*, 831–8.
- 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.
- Laruelle, M. (2000). The role of endogenous sensitization in the pathophysiology of schizophrenia: implications from recent brain imaging studies. *Brain Research - Brain Research Reviews*, *31*(2–3), 371–84.
- Miikkulainen, R. (1993). *Subsymbolic natural language processing: An integrated model of scripts, lexicon, and memory*. Cambridge, MA: MIT Press.
- Pollack, J. (1990). Recursive distributed representations. *Artif Intell*, *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.
- Spitzer, M. (1997). A cognitive neuroscience view of schizophrenic thought disorder. *Schizophrenia Bull*, *23*(1), 29–50.