Appeared in: B. C .Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society* (pp- 2083-2088). Austin, TX: Cognitive Science Society

Activation or Inhibition? Why Reasoners are Not Blind for Alternative Explanations

Katja Mehlhorn (katja.mehlhorn@phil.tu-chemnitz.de)*
Martin R. K. Baumann (martin.baumann@dlr.de)**
Franziska Bocklisch (franziska.bocklisch@phil.tu-chemnitz.de)*
*Chemnitz University of Technology, Department of Psychology,
Wilhelm-Raabe-Str. 43, 09107 Chemnitz, Germany
**German Aerospace Center, Lilienthalplatz 7, 38108 Braunschweig, Germany

Abstract

Reasoners can reduce the complexity of diagnostic reasoning tasks by limiting the search for an observation's explanation to those candidate explanations that show the highest consistency to previous observations. This strategy might, however, impair the availability of explanations beyond that limited set of explanations. In two experiments, where reasoners needed to reconsider such excluded alternatives for explaining inconsistent data in a well practiced diagnostic reasoning task, we found evidence that this is not necessarily the case. By tracking the availability of different explanations with a probe reaction task, we were furthermore able to show how availability seems to be regulated by automatic activation, rather than inhibition processes.

Keywords: abductive reasoning; hypothesis generation; inconsistent data; anomaly; automatic processes; activation; inhibition

Introduction

Generating explanations for one or more observations is a key component of many real-world tasks, such as medical diagnosis, software debugging, scientific discovery, and social attribution. This kind of reasoning is often called abductive reasoning and described as highly complex because of the amount of explanations possibly accounting for each observation (e.g. Bylander, Allemang, Tanner, & Josephson, 1991). One commonly proposed mechanism to reduce this complexity is to consider not all, but only a subset of all possible explanations for an observation (Johnson & Krems, 2001; Josephson & Josephson, 1994; Thomas, Dougherty, Sprenger, & Harbison, 2008). However, by adopting this strategy, reasoners may become unaware of alternative, potentially better explanations. Encountering an observation that can only be accounted for by such an overlooked alternative would then represent an anomaly.

Although several theoretical models related to abductive reasoning in general have been proposed (e.g. Arocha & Patel, 1995; Johnson & Krems, 2001; Thagard, 1991; Thomas et al., 2008), only little research was directed at how reasoners deal with anomalous or inconsistent data. Which cognitive processes allow this performance despite the limitations of human working memory? Based on general complexity-reducing mechanisms postulated by previous research, we discuss how automatic memory activation processes can moderate the availability of alternative explanations when an anomaly is encountered.

With increasing experience in a domain, reasoners develop knowledge structures that reflect the structure of the task. For example, a physician learns, with the increasing number of patients encountered, which symptoms are associated with which diseases and how frequently the association occurs. This frequency of co-occurrence is reflected in the reasoner's knowledge structure by the strength of association between symptoms and diseases causing them (Kintsch, 1998; Thomas et al. 2008). Given such an adapted knowledge structure, data extracted from the environment can serve as a cue for the retrieval of diagnostic hypotheses from long-term memory.

Applying Kintsch's (1998) construction integration theory to abductive reasoning, Baumann et al. (2007b) provided a model of how implicit, automatic memory activation processes enable people to limit the search for possible explanations to those that are most plausible in the current context. When observing a new symptom, explanations highly associated with that symptom are activated from memory. If any further information is available (e.g. possible explanations for previous symptoms), the activated explanations are integrated with that information via spreading activation and inhibition. Integration leaves those explanations highest activated that show the highest consistency to all encountered observations. Thus, at least on the first run, only those potential explanations are considered, that turned out to be highly relevant in the past and that are also compatible to the current explanatory context (see Thagard (1991) and Thomas et al. (2008) for similar assumptions).

But what happens if a later observation can be explained only by an explanation that was previously excluded from consideration due to this mechanism? We propose that, provided high domain expertise, automatic activation processes also moderate the availability of previously excluded alternative explanations. The literature, however, does not allow clear predictions about the nature of these activation processes.

First, one could assume that automatic processes facilitate switching between explanations. Depending on the structure of the domain knowledge, explanations in memory are interlinked. Thus, after observing a new symptom, spreading activation should not only lead to the activation of explanations directly linked to the symptom, but also to the activation of other, related explanations. Therefore, one could then assume that alternative explanations are easily reconsidered, if they are related to considered explanations in some way.

Alternatively, one could assume that after an explanation is rejected and removed from consideration, it is actively inhibited to minimize interference. Such results are often reported in discourse or text comprehension research. For example, May, Zacks, Hasher, and Multhaup (1999) showed that successfully reanalysing garden path sentences highly depends on people's abilities to inhibit no longer appropriate interpretations. Spreading inhibition processes could then result in the inhibition not only of the rejected explanation, but of other, related explanations. Following such an inhibition pattern, it should be difficult to consider an alternative hypothesis that is somehow related to the to-be-rejected hypotheses forming the current explanation.

Experiment 1

Our aim in this experiment was to test whether the strategy of limiting the search for explanations to those being consistent with all previous observations indeed impairs the availability of alternative explanations. Therefore, we compared participants' performance when diagnosing explanations for either consistent or inconsistent sets of symptoms presented sequentially. During trials, we tracked the activation of different explanations with a probe reaction task to test the different assumptions about how automatic memory activation processes moderate the availability of explanations.

Material and Predictions

Given the importance of domain knowledge and therewith highly adapted retrieval structures regulating the knowledge's availability, precise measurement of the availability of different explanations requires a strict control of the domain knowledge and retrieval structures. To allow for this control, we used an artificial diagnosis task, the "chemical accident" task, rather than real world knowledge (Baumann et al, 2007 a and b). In this task, participants are presented with the symptoms of hypothetical patients after a chemical accident and have to identify the chemicals that caused these symptoms. To be able to solve this task, participants first have to learn about the chemicals and the symptoms they cause. The knowledge base used in this experiment consisted of nine chemicals that were grouped into three hypothetical categories (Table 1).

Table 1: Domain knowledge participants had to acquire before experiment 1 (original material in German).

Category	Chemical	Symptoms
Landin	B	cough, short breath, headache, eye inflammation
	T	cough, short breath, headache, itching
	W	cough, eye inflammation, itching
Amid	Q	skin irritation, redness, headache, eye inflammation
	M	skin irritation, redness, headache, itching
	G	skin irritation, eye inflammation, itching
Fenton	K	diarrhea, vomiting, headache, eye inflammation
	H	diarrhea, vomiting, headache, itching
	P	diarrhea, eye inflammation, itching

Each chemical caused three or four symptoms that themselfes appeared either in one category (specific symptoms, e.g. cough), or in all categories (unspecific symptoms, e.g. headache). This hierarchical structure was supposed to ease the learning of the material and reflects the hierarchical knowledge organization found in medical diagnosis in simplified form (Arocha & Patel, 1995).

After acquiring the knowledge in a pre-experimental training session, participants solved various trials. In each trial, symptoms of one patient were presented sequentially (each for 2 seconds with a 1 second ISI) before the participant was asked for the chemical that caused these symptoms. During trials, symptoms could appear either in a strong form (indicated by bold letters) or in a weak form (indicated by normal letters). Whereas strong symptoms were always caused by a chemical, weak symptoms could also be induced by other reasons (e.g. short breath by physical arousal). Due to this manipulation, we were able to create *inconsistent trials* by inserting an additional misleading weak symptom into the trial. Parallel to each inconsistent trial, we created *consistent trials* in which the same weak symptom was presented but without being misleading (see Figure 1 for examples)[1].

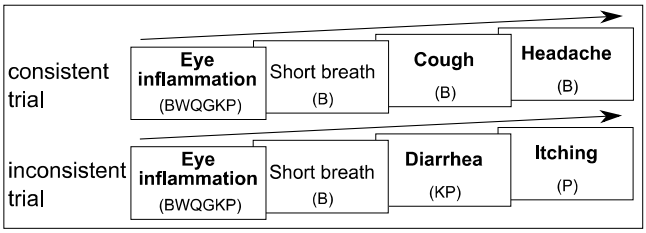

Figure 1: Examples for a consistent and an inconsistent trial in Experiment 1. Letters in parentheses represent possible explanations for all symptoms presented up to this point.

In both types of trials, the initial (strong) symptom was linked to various explanations (BWQGKP in the example trials). The second (weak) symptom was consistent with only one of these explanations and thereby allowed for reducing the number of considered explanations (to chemical B in the example). Symptom three (strong) was the critical symptom. It either continued to be consistent with the previous explanations (in consistent trials) or it was inconsistent with this explanation (in inconsistent trials). To solve inconsistent trials, participants needed to realise that in this case the weak symptom was not caused by a chemical but by some other reason. Thus, the chemicals that had been rejected due to this symptom needed to be reconsidered and some of them were able to explain the critical third symptom (K and P in the example). The fourth (strong) symptom in both types of trials supported an explanation consistent with the third symptom.

Depending on how automatic activation or inhibition processes influence performance in such a task one would expect different levels of activation for the different explanations in inconsistent trials. For explanations that are

[1] Presenting symptoms in weak form also in consistent trials made sure that participants used these symptoms to generate explanations in the same way as they used strong symptoms.

considered as relevant after the first and second symptom but need to be rejected after observing the inconsistent third symptom (*rejected after 3rd symptom*), both assumptions (activation and inhibition) would predict a decrease in activation after the third symptom. This decrease should, however, be more accentuated and even go below baseline level if rejected explanations are inhibited. For explanations that were misleadingly rejected after the second and need to be reconsidered after the third symptom (*reconsidered after 3rd symptom*) one would expect increased activation after the third symptom. If these explanations would benefit from their pre-activation by the current explanation, this decrease should, however, be much more accentuated than if they received inhibition from the rejected current explanation.

To test these assumptions we tracked the explanations' activation with a probe reaction task. After one of the symptoms in each trial, a probe was presented and participants had to decide as fast as possible whether or not this probe was the name of a chemical. Comparable to lexical decision tasks, one would expect the response times for this decision to be the shorter, the higher available a chemical currently is as explanation in working memory. Analogous to neutral words in lexical decisions tasks, we used the names of chemicals that are irrelevant as explanations for the presented symptoms to assess explanations' baseline activation.

Participants, Procedure, and Design

Twenty-two undergraduate students (14 female and 8 male; mean age 22.77, $SD = 3.57$) from Chemnitz University of Technology took part in this experiment that consisted of one session for training and two test sessions. In the training session participants acquired the domain knowledge and practiced the chemical accident task until achieving a level of at least 80% correct trials. The test sessions each started with a practice block to refresh the knowledge and then participants solved 144 test trials.

In each test session, one third of these trials were *inconsistent* and one third were *consistent* as described above. In consistent trials, we measured the activation of hypotheses before the critical third symptom. This allowed us to use inconsistent trials exclusively to measure hypotheses' activation after the inconsistency became apparent, and thus, keep the proportion of inconsistent trials as low as possible. The remaining third were filler trials with different symptom order, also consisting of weak and strong symptoms. Trials were presented in randomized order.

Probes were presented after one of the symptoms in each trial and were to the same percentage either targets (the names of the 9 chemicals – that were at point of measurement considered, rejected, or irrelevant as explanations), or distractors (9 different letters). The number of symptoms before presentation of the probe as well as the type of probe were balanced across trials and presented in randomized order.

Dependent variables were accuracy of the final diagnoses and reaction time of correct probe reactions.

Results

Diagnosis performance in consistent and inconsistent trials. In inconsistent trials, participants solved 88.6% of the trials correctly. Although this performance is considerably high, it differs significantly from the performance in consistent trials that was 93.2%, $t(21) = -2.65$, $p = .015$.

Time course of hypotheses activation during inconsistent trials. Figure 2 shows reaction times to targets probing the three courses of explanations between the second and the fourth symptom.

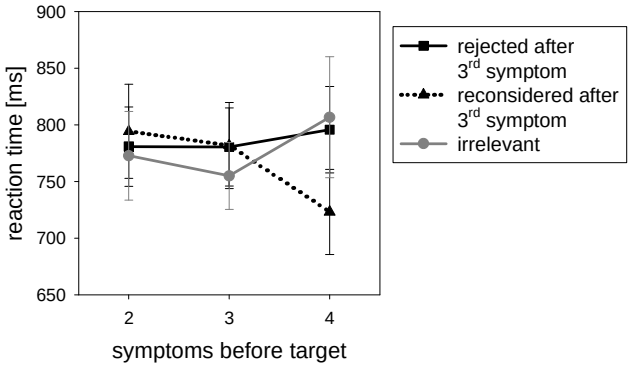

Figure 2: Reaction times to targets probing the activation of different courses of explanations during inconsistent trials in Experiment 1.

After the second symptom, considered targets (solid) seemed to be responded to slightly faster than misleadingly rejected ones (dotted). After symptom 3, to-be-rejected (solid) and to-be-reconsidered (dotted) targets produced the same reaction time. After symptom four, the postulated effect of switching between explanations became visible. Reconsidered targets now highly increased in reaction speed, whereas reaction speed for targets rejected after the third symptom decreased. Interaction of these two courses was analyzed with a repeated measures ANOVA (Type of target (*rejected after 3rd symptom* and *reconsidered after 3rd symptom*) vs. symptoms before target (*2,3, or 4*)) and was significant, $F(2.94,61.69) = 3.30$, $p = .016$, $\eta^2 = .18$.

Targets labeled irrelevant were intended to be a measure for baseline activation of explanations. However, data revealed that reaction times for theses targets did not follow the expected pattern. Reaction times being not only below those for rejected but also below those for relevant targets after the second and third symptom indicate that these targets did not measure the baseline activation of irrelevant explanations. A possible explanation can be found in the structure of the presented trials. Because all trials in this experiment started with unspecific symptoms (each linked to six out of the nine possible chemicals), it seems reasonable that after the first symptom, participants considered all chemicals as somehow relevant and waited for the second symptom to limit the number of considered explanations. Hence, irrelevant targets might not have tracked the activation of irrelevant explanations, but rather that of explanations rejected after the second symptom.

Discussion

As expected, participants solved less inconsistent trials correctly than consistent trials. Results of the probe task reveal that as predicted, after the critical third symptom, reconsidered explanations regained activation and explanations that had to be rejected lost activation. This interaction did however not show up directly after the third symptom, but only after the fourth symptom. Thus, it seems that participants had some difficulty reconsidering previously rejected explanations directly after the inconsistency occurred, but could use additional evidence to solve that anomaly. These results match findings of Keinath and Krems (1998) who reported that anomaly resolution was facilitated by an increasing amount of evidence supporting the anomaly. Unfortunately, we failed to assess baseline activation in this experiment and therefore cannot draw clear conclusions about the existence of inhibition processes. The small decrease of explanations' activation after their rejection, however, suggests that after being rejected, explanations' activation simply decays rather than being inhibited. The fact that participants finally managed to reconsider the previously rejected alternative in most inconsistent trials furthermore contradicts the inhibition assumption.

Concluding, although participants diagnosis performance was lower in inconsistent than in consistent trials, the still high performance and the regained activation of previously rejected alternatives indicates that limiting the number of explanations considered does not necessarily blind the reasoner to possible alternative explanations. However, the results do not allow to decide clearly between the different assumptions about the nature of activation processes underlying this performance.

To be able to distinguish between these two alternatives, we designed an experiment where the proposed effects of either mechanism should be even stronger. According to definition, spreading activation, as well as spreading inhibition processes should influence the activation level of other explanations the more the closer explanations are related to each other. Considering the hierarchical organization of knowledge in many domains, one aspect that determines explanations' relatedness is whether they belong to the same explanatory category or to different categories.

Baumann, Bocklisch, Mehlhorn, and Krems (2007a) report that reasoners had more difficulties in solving anomalous trials when they required switching between explanations within one category, than in solving anomalies requiring switching to an explanation from a different category. This result contradicts previous results on anomaly resolution (e.g. Klahr & Dunbar, 1988). It could, however, possibly be explained by inhibition processes. If an explanation is rejected and removed from consideration by inhibition processes, spreading inhibition could lead to the inhibition of other, closely related explanations. Thus, when the reasoner needs to switch between explanations that are closely related, rejection of one explanation should reduce the availability of alternative explanations from the same category. Spreading activation processes could in contrast easily explain results as reported by Klahr and Dunbar. If unconsidered explanations receive activation from the current explanation, the amount of received activation should be the higher the closer explanations are interlinked. Thus, switching between explanations from one category should be easier than switching between different categories. To distinguish between these assumptions, we conducted a second experiment in which inconsistent trials required participants to reconsider explanations from the same category as the current explanation.

Experiment 2

The goal of the second experiment was to examine diagnosis performance and hypothesis activation, when the reasoner has to switch between explanations that are more closely related to each other than explanations were in Experiment 1. For this purpose, we constructed trials where the reasoner had to switch from one explanation to another within the same category. Figure 3 illustrates such a case. The initial symptom of the presented trial (cough) can be caused by three chemicals from one category (BTW). The second symptom (vomiting) allows for limiting the number of considered explanations to one chemical from this category (T). In the consistent version of the trial, symptoms three and four confirm this chemical as explanation. In the inconsistent version, symptoms three and four contradict this explanation and require the reasoner to switch the explanation from chemical T to chemical B from the same category.

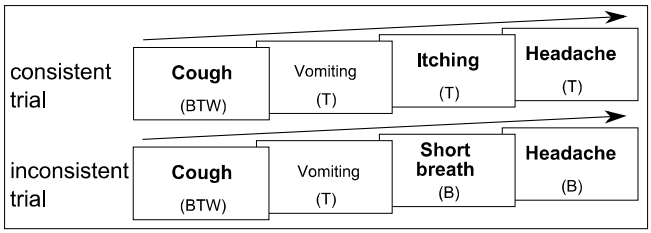

Figure 3: Examples for a consistent and an inconsistent trial in Experiment 2. Letters in parentheses represent possible explanations for all symptoms presented up to this point.

Material and Predictions

For this experiment, we used the material of Experiment 1, with minor adaptations to allow for generating inconsistent trials requiring a hypothesis change within categories. Predictions for hypothesis activation and participants' diagnosis performance were the same as in Experiment 1, but effects should be greater now.

Participants, Procedure, and Design

Twenty-three undergraduate students (13 female and 10 male; mean age 24.04, $SD = 6.66$) from Chemnitz University of Technology, who had not participated in Experiment 1, took part in this experiment. After acquiring the domain knowledge and practicing the chemical accident task in a pre-experimental training session, they completed 4 experimental test sessions of 96 trials each.

In comparison to Experiment 1, the proportion of each type of trial was lowered to allow for using more filler trials. Half of all trials were filler trials, 24 starting with specific symptoms (linked to explanations from only one category) and 24 starting with unspecific symptoms (linked to explanations form various categories). This should help to avoid the problem that participants did not exclude irrelevant explanations from consideration after the first symptom as in Experiment 1. Half of all trials were test trials, of which 24 were *inconsistent* and 24 were *consistent* trials. Consistent trials again resembled inconsistent trials until the second symptom but then ended without an anomaly.

As in Experiment 1, diagnosis performance was measured by the accuracy of the final diagnoses and hypothesis activation was measured with the probe reaction task. Important factors (type of trial, type of probe, and number of symptoms presented before probe) were balanced and randomized.

Results

Diagnosis performance in consistent and inconsistent trials. When switching within categories was necessary to solve the contradiction in an inconsistent trial, participants solved 96.1% of the consistent and 94.8% of the inconsistent trials correctly.

Time course of hypotheses activation during inconsistent trials. Figure 4 shows reaction times to targets probing the three courses of explanations between the second and the fourth symptom.

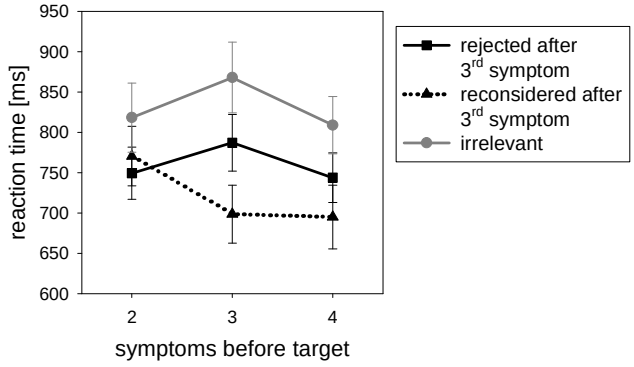

Figure 4: Reaction times to targets probing the activation of different courses of explanations during inconsistent trials in Experiment 2.

As in Experiment 1, after the second symptom, misleadingly rejected explanations produced slightly slower response times than considered explanations. After the critical third symptom, we now, however, observed faster reaction times for to-be-reconsidered explanations and slower times for to-be-rejected explanations. After the fourth symptom, reaction time differences between these two types of explanations remained, although rejected explanations were responded to faster than after the third symptom. This interaction of the factors type of target (*rejected after 3rd symptom, reconsidered after 3rd symptom, or irrelevant*) and symptoms before target (*2,3, or 4*) was significant, $F(3.34,73.19) = 4.36$, $p = .005$, $\eta^2 = .17$. We furthermore observed a main effect for the factor type of trial. After each symptom, irrelevant explanations produced the slowest reaction times, $F(2,44) = 28.88$, $p < .001$, $\eta^2 = .57$.

Discussion

The high diagnosis accuracy for inconsistent trials in this experiment shows that participants did not seem to have difficulties switching between explanations within one category. This conclusion is supported by the results of the probe reaction task.

After the critical third symptom, reconsidered explanations regained activation quickly. The fourth symptom additionally supporting the reconsidered explanation did not lead to a further increase, indicating that this additional support was not necessary to solve the anomaly. Thus, misleadingly rejected explanations from the same category as the current explanation could be reconsidered easily in this experiment - a result supporting the assumption of spreading activation and contradicting that of inhibition.

Explanations that had to be rejected after the critical third symptom did lose activation but stayed far above the baseline level of irrelevant explanations. Referring to literature on text comprehension, we assumed that inhibition of explanations should result in an activation level below baseline. This assumption was not supported by our results. However, one still could argue that inhibition first would have to overcome the previous activation before driving RTs above baseline, which might or might not occur.

Concluding, although results of this experiment do not completely contradict inhibition, they are more compatible with an activation-only account.

Conclusions

A commonly proposed mechanism to reduce complexity in abductive reasoning tasks is to consider not all possible explanations for an observation but only those consistent with previous observations. This strategy might, however, impair the reasoners awareness for alternative explanations that were thereby excluded from consideration. To test this assumption, we conducted two experiments where participants needed to reconsider such excluded alternatives for explaining inconsistent data. In the experiments, we tracked the availability of explanations with a probe reaction task to distinguish between two possibilities of how automatic memory activation could moderate the availability of such explanations.

In both experiments, participants showed a high diagnostic performance also in inconsistent trials. This suggests that in well practiced tasks reasoners do not necessarily have difficulties reconsidering explanations that they excluded from consideration before. When these

explanations stem from another category than the current explanations (as in Experiment 1), this reconsideration seems, however, more difficult than when the explanations stem from the same category (as in Experiment 2).

The results of the probe reaction tasks support this finding. In both experiments previously rejected explanations regained activation when they needed to be reconsidered to account for inconsistent data. This regain happened more promptly when explanations from the same category as the current explanation were to be reconsidered (as in Experiment 2), than when participants needed to reconsider explanations from another category (as in Experiment 1). Interpreting the results of such a direct comparison between two different experiments needs to be done with care. We therefore plan to conduct future research combining the two conditions (switching between and within categories) in one experiment. However, our results already provide some evidence for the assumption that to-be-reconsidered explanations benefit from the pre-activation of their category by the current explanation and that this pre-activation is the stronger the closer related explanations are in memory. Thus, switching between closely related explanations seems to be easier because such explanations are higher available to the reasoner than explanations from unconsidered categories.

The conclusions we can draw about possible inhibition processes are less clear. In Experiment 1, we failed to measure the activation of really irrelevant explanations and thus, we cannot compare any activation level to a baseline. The results of Experiment 2 do not support inhibition, but since the inhibition just might not have been strong enough to overcome the previous activation, they also do not clearly contradict it. Moreover, the time delay (SOA) between symptoms and probes might have been inadequate to find inhibition. Ever since the early findings of Neely (1977), it has been known that activatory and inhibitory processes exhibit different time courses. Which SOAs are applicable to find possible inhibition effects in abductive tasks will have to be an interesting question for further research.

Our results also shed some light on the findings of Baumann et al (2007a). Contrary to other researchers, they found a higher performance when switching between categories than when switching within categories was necessary to explain inconsistent data. In their experiment, the critical symptoms inducing the anomaly varied in their specifity. Symptoms requiring a reconsideration of explanations from the same category were always unspecific and thus linked to explanations from several categories. Such symptoms might be less able to activate linked explanations than symptoms that are linked to explanations from only one category. Thus, the novel findings of Baumann et al. might be moderated by the symptom characteristics influencing the activation processes.

To conclude, our findings show that the strategy of reducing the complexity of abductive tasks by limiting the search for explanations to those consistent with all previous observations must not necessarily impair the availability of possible alternative explanations. We could furthermore provide evidence that in well practiced tasks this effect might be due to automatic spreading activation processes.

Acknowledgements

We thank Georg Jahn, Josef Krems and four anonymous reviewers for their helpful comments on earlier drafts of this paper.

References

Arocha, J. F., & Patel, V. L. (1995). Construction-integration theory and clinical reasoning. In Ch. A. Weaver, S. Mannes, & C. R. Fletcher (Eds.), *Discourse comprehension. Essays in honor of Walter Kintsch* (pp. 359-381). Hillsdale, NJ: Lawrence Erlbaum Associates.

Baumann, M. R. K., Bocklisch, F., Mehlhorn, K., & Krems, J. F. (2007a). Changing explanations in the face of anomalous data in abductive reasoning. In D. S. McNamara & J. G. Trafton (Eds.), *Proceedings of the 29th Annual Cognitive Science Society* (pp. 797-802). Austin, TX: Cognitive Science Society.

Baumann, M. R. K., Mehlhorn, K., & Bocklisch, F. (2007b). The activation of hypotheses during abductive reasoning. In D. S. McNamara & J. G. Trafton (Eds.), *Proceedings of the 29th Annual Cognitive Science Society* (pp. 803-808). Austin, TX: Cognitive Science Society.

Bylander, T., Allemang, D., Tanner, M. C., & Josephson, J. R. (1991). The computational complexity of abduction. *Artificial Intelligence, 49(1-3)*, 25-60.

Johnson, T. R., & Krems, J. F. (2001). Use of current explanations in multicausal abductive reasoning. *Cognitive Science, 25(6)*, 903-939.

Josephson, J., & Josephson, S. G. (1994). *Abductive inference: Computation, philosophy, technology*. New York, NY: Cambridge University Press.

Keinath, A. & Krems, J. F. (1998). The influence of anomalous data on solving human abductive tasks. In L. Magnani, N. Neressian, & P. Thagard, *Philosophica, Abduction and Scientific Discovery [Special Issue], 61(1)*, 39-50.

Kintsch, W. (1998). *Comprehension: A paradigm for cognition*. New York, NY: Cambridge University Press.

Klahr, D., & Dunbar, K. (1988). Dual space search during scientific reasoning. *Cognitive Science, 12*, 1-48.

May, C. P., Zacks, R. T., Hasher, L., Multhaup, K.S., (1999). Inhibition in the processing of garden path sentences. *Psychology and Aging. 14*, 304–313.

Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General, 106*, 226-254.

Thomas, R. P., Dougherty, M. R., Sprenger, A. M., & Harbison, J. I. (2008). Diagnostic hypothesis generation and human judgment. *Psychological Review, 115(1)*, 155-185.

Thagard, P. (1991). Defending explanatory coherence. *Behavioral and Brain Sciences, 14*, 745–748.