Making the Implausible Plausible

Louise Connell (louise.connell@ucd.ie)
Department of Computer Science, University College Dublin
Belfield, Dublin 4, Ireland

Abstract

Whenever we must evaluate a theory, consider an excuse, or appraise a situation, we must judge how plausible things appear to us. In short, plausibility judgement occupies a central position in human cognitive life. Recently, it has been shown that the plausibility of a scenario depends on how its events can be connected (Connell & Keane, in press). In this paper, two experiments examine how a normally implausible scenario can be made to seem plausible by forcing a connection between its events. Results show that people's perceptions of a scenario's plausibility can be manipulated by encouraging them to represent events in a causal chain or temporal sequence.

Introduction

Every day, in many different situations, we judge plausibility. Whether evaluating a theory, considering the plot quality of a movie, or listening to child explain how a dish came to be broken, we are assessing how plausible a scenario seems to us.

Across the cognitive science and cognitive psychology literature, plausibility judgement has been shown to be useful in a diverse range of cognitive tasks. People often use plausibility judgements in place of costly retrieval from long-term memory, especially when verbatim memory has faded (Lemaire & Fayol, 1995; Reder, 1982; Reder, Wible & Martin, 1986). Plausibility is also used as a kind of cognitive shortcut in reading, to speed parsing and resolve ambiguities (Pickering & Traxler, 1998; Speer & Clifton, 1998). In everyday thinking, plausible reasoning that uses prior knowledge appears to be commonplace (Collins & Michalski, 1989), and can even aid people in making inductive inferences about familiar topics (Smith, Shafir & Osherson, 1993). It has also been argued that plausibility plays a fundamental role in understanding novel word combinations by helping to constrain the interpretations produced (Costello & Keane, 2000; Lynott, Tagalakis & Keane, in press). Yet, despite its apparent usefulness in cognitive life, the study of plausibility judgement in its own right has been neglected in cognitive science until recently.

The Knowledge-Fitting Theory of Plausibility

Recently, Connell and Keane have proposed the Knowledge-Fitting Theory of Plausibility (2003a, 2003b, in prep; Connell, 2004) to rectify this oversight. According to the Knowledge-Fitting Theory, a plausible scenario is one that fits well with our knowledge of the world. In other words, a plausible scenario has good *concept-coherence*. A concept-coherence view of plausibility suggests that when people make a plausibility judgement, they relate the current scenario to their prior experience, and in some way assess whether it fits in with what they have experienced in the past. For example, take the scenario "*The bottle rolled off the shelf and smashed on the floor.*" People might understand this scenario by drawing a causal inference between the events – that is, the bottle falling *caused* it to smash on the floor. This might lead them to judge this scenario as being highly plausible because prior experience tells them that fragile things often break when they fall on hard surfaces. Put simply, the scenario has a certain concept-coherence. In contrast, if the scenario was "*The bottle rolled off the shelf and melted on the floor*", people might consider it less plausible because they cannot connect the events, as their past experience has few examples of falling fragile objects melting on contact with floors. In other words, this scenario lacks a certain concept-coherence.

The Knowledge-Fitting Theory is supported by a number of empirical findings. For example, Black, Freeman and Johnson-Laird (1986) disrupted the sequence of events in short stories and found that people's plausibility ratings were sensitive to the degree to which the overall concept-coherence of the story had been altered; when people could no longer infer connections between events, they no longer considered the stories plausible. Connell and Keane (in press; Connell, 2004) investigated this issue further, and found that the plausibility of a scenario is not only affected by *whether* the events can be connected, but also by *how* the events are connected. For example, events linked by causal inferences (i.e., *X* was caused by *Y*) were judged to be more plausible than events linked by temporal inferences (i.e., *X* happens after *Y*). These types of scenario were both considered more plausible than scenarios where the events could not be connected at all (i.e., unrelated events). Connell and Keane (2003a, 2003b, in prep.) suggest that causal connections have better concept-coherence than temporal connections because they fit more closely with prior experience, and this makes causal scenarios seem more plausible.

The Current Study

When we judge the plausibility of some scenario in everyday life, it is often with the objective of accepting or rejecting the presented scenario. For example, we may choose to accept or reject an excuse based on whether we find it plausible or not. However, plausibility is not always a binary variable (i.e., a choice between plausible and implausible) (e.g., Black et al., 1986; Connell & Keane, in press). Rather, it may sometimes be considered as a sliding scale between plausibility and implausibility, and we may judge a particular scenario as lying somewhere along this scale. In addition, the plausibility of a particular scenario may not have a constant value: for example, temporal scenarios can be made to seem less plausible when attention is drawn to their non-causal nature (Keane, Connell & O'Donoghue, in prep.).

This paper investigates whether an implausible scenario can be made seem plausible by forcing a particular connection between its events. The Knowledge-Fitting Theory holds that if we cannot connect the events in scenario, we will find it implausible. However, it is possible that if we manage to connect unrelated events in some way, then the plausibility of the scenario might increase. For example, the scenario "*The bottle rolled off the shelf and melted on the floor*" seems generally implausible, but it is possible to construct some set of circumstances that makes it appear more plausible (e.g., the bottle melted because the floor was very hot, because the house was on fire). This suggests that the concept-coherence and plausibility of a normally implausible scenario could be manipulated by encouraging people to make particular connections between events. Therefore, the first experiment asks people to make specific causal or temporal connections between unrelated events, and examines how this influences their decision to accept or reject the scenario. In the second experiment, the same manipulation is used to show how different connections also influence people's plausibility ratings. These results are then related back to the Knowledge-Fitting Theory, and are used to examine what kind of relationship exists between binary and scale plausibility judgements.

Experiment 1

In Experiment 1, participants are presented with implausible scenarios and are asked to judge whether each scenario is plausible or not. Connell and Keane (in press) show that scenarios with no connection between events are considered implausible, while those with causal and temporal connections are considered plausible. This experiment leads people to think about how the events in a scenario may be causally or temporally connected, and examines how these different connections can make people accept as plausible a scenario that would normally be rejected as implausible. For example, take this scenario: "*The teacher misspelled a word. The vase smashed.*" If people are encouraged to represent this normally implausible scenario with a causal or temporal connection between events, then this may lead them to perceive the scenario as plausible. For example, if this scenario was represented within a specific temporal frame (e.g., the vase smashed a second or two after the teacher misspelled a word), then it may have sufficient concept-coherence to appear plausible to some people. Alternatively, if this scenario was represented with a causal chain between events (e.g., the vase smashed because the teacher bumped against it when taking a step back to examine the misspelled word), then it may have sufficient concept-coherence to appear plausible to many people.

This experiment uses different types of question to encourage people to make particular connections between events. In the no-connection control condition, participants judged the plausibility of the scenario directly after reading it, and were not asked to connect the events in any particular way. In the other conditions, participants had to answer a question that encouraged them to represent the scenario in a particular way before judging its plausibility; the causal condition required participants to connect the events as cause and effect, while the temporal condition required participants to connect the events as a temporal sequence. In short, people are expected to find few scenarios plausible in the no-connection condition (poor concept-coherence), more scenarios plausible in the temporal condition (better concept-coherence), and most scenarios plausible in the causal condition (best concept-coherence).

Method

Materials & Design. Materials consisted of twenty "implausible" scenarios, each consisting of two sentences describing unrelated events. The materials were constructed by creating twenty causal scenarios (e.g., "*The surgeon performed the operation. The patient recovered.*") and then randomizing the combinations of first and second sentences (e.g., "*The surgeon performed the operation. The candle flickered.*"). Thus, each scenario contained two events where the cause (Event A) was followed by a different, unrelated effect (Event B). The experimental design was a single between-participants factor (connection type), with three conditions (causal, temporal, no-connection). A between-participants design was chosen to avoid possible confounds (e.g., participants forming a causal connections in all presented scenarios).

Participants. Thirty-six student volunteers from University College Dublin participated in this experiment.

Procedure. Participants were randomly assigned, in equal numbers, to one of the conditions in the experiment.

Instructions stated that each scenario was taken from a story and consisted of two events: Event A and Event B. An example not used in the materials was given:

Event A: *The boy kicked the football.*

Event B: *The branch snapped.*

In the causal condition, participants were asked to write down their answer to the question "Why do you think Event B happened?" and were presented with a sample answer "The branch snapped because the football hit it hard, because the boy was aiming at the tree." In the temporal condition, participants were asked to write down their answer to the question "How long after Event A do you think Event B happened?" and were presented with a sample answer "The branch snapped 2 or 3 seconds later." In all conditions, participants were then given a forced-choice plausibility judgement "Do you find this scenario plausible?" and asked to circle "Yes" if they would accept the scenario as plausible, and to circle "No' if they would reject the scenario as implausible. The scenarios were presented on separate pages in random order, resampled for each participant.

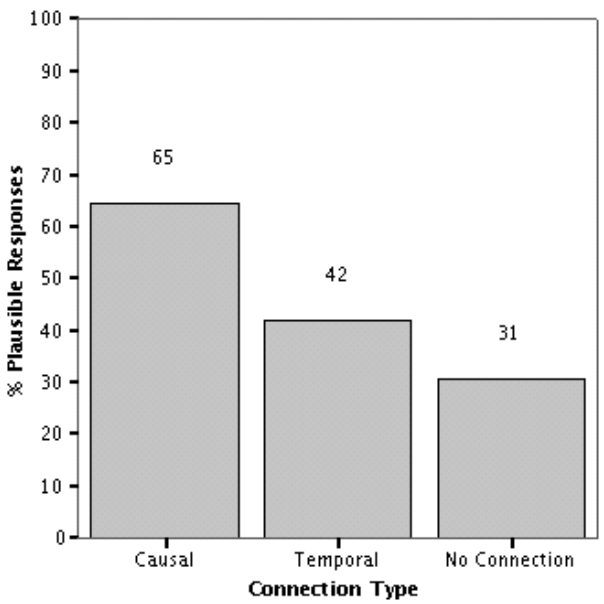

Figure 1: Percentage of scenarios accepted as plausible for each connection type in Experiment 1.

Results & Discussion

The results were in line with predictions, and are shown in Figure 1. Three scenarios were considered plausible by more than 80% of participants in the no-connection control condition, and these were excluded from the analysis in all conditions. While this paper lacks sufficient space to discuss participants' answers in detail, people reported little difficulty in making the connections in the temporal condition, and only occasional difficulty in the causal condition.

People's willingness to accept the scenarios as plausible was influenced by how they had been encouraged to represent the connections between events, as shown by chi-squared analysis, $\chi^2 = 47.74$, $df = 2$, $p < 0.0001$. When asked to represent the events as a specific temporal sequence, people accepted significantly more scenarios as plausible (42%, compared to 31% control), $\chi^2 = 5.31$, $df = 1$, $p < 0.05$. However, the greatest change occurred when people were asked to represent a causal chain between the events, with over twice the number of scenarios being perceived as plausible (65%, compared to 31% control), $\chi^2 = 45.77$, $df = 1$, $p < 0.0001$. In this way, causal connections were reliably better than temporal connections at making scenarios appear plausible (65% compared to 42%), $\chi^2 = 20.71$, $df = 1$, $p < 0.0001$.

So what makes implausible scenarios suddenly appear plausible? Why do people perceive the plausibility of a scenario differently when they are asked why or when events happened? After all, the actual connection made between unrelated events is arbitrary: people are free to come up with any explanation or time frame they choose to connect the events. There is nothing stopping people from causally connecting the events in the no-connection control condition, nor is anything stopping them from causally representing the scenario when answering the temporal question. Indeed, it could be argued that the 31% plausible responses in the no-connection control condition result from the times that people managed to make a causal connection between events without any guidance. In effect, the pattern of results suggests that a kind of "cognitive laziness" is at play, and that people do not put any more effort into representing the scenarios than is absolutely necessary. In the no-connection control condition, most scenarios are judged as implausible because no obvious connection between the events comes to mind. However, in order to be able to answer the temporal question, a certain amount of extra effort must go into connecting the events. If the resulting representation has sufficient concept-coherence, the scenario then seems plausible. Lastly, answering the causal question requires quite a lot of effort, as people must explicitly lay out the circumstances that brought about the second event. The resulting causal representation is likely to have good concept-coherence, and so the scenario is likely to appear plausible. This "cognitive laziness" view is consistent with other studies that have demonstrated people's reluctance to infer causal relations unless prompted to do so (e.g., Keane, 1997; McKoon & Ratcliff, 1992).

This experiment shows that people will accept an implausible scenario as plausible if they are encouraged to connect events in a certain way. However, it is also possible to elicit a more fine-grained judgement of plausibility. When people accept a scenario, is it because they judge it to be very plausible or just moderately plausible? This question is addressed in the next experiment.

Experiment 2

Experiment 2 is identical to Experiment 1, except participants are asked to rate the plausibility of scenarios on a scale from 0-10 instead of simply choosing whether the scenario seems plausible or implausible. In other words, participants are presented with implausible scenarios and are asked to judge how plausible they find each scenario. As with Experiment 1, scenarios in the causal and temporal conditions are expected to be rated as more plausible than those in the no-connection control condition because of their greater concept-coherence. However, there are two possibilities for how the causal and temporal conditions may be distinguished.

The first possibility is that when people in the causal condition accept a scenario as plausible, they actually consider it to be *highly* plausible. This means that causal connections between unrelated events would be considered to have very good concept-coherence; indeed, just as good as for more straightforward causal scenarios. If this were the case, the results would mirror those of Connell and Keane (in press, experiment 1), where causal scenarios were rated as highly plausible (7.8 out of 10) and temporal scenarios were rated as only moderately plausible (4.2 out of 10).

The second possibility is that, although implausible scenarios may become acceptably plausible in the causal condition, they will never seem *highly* plausible. This means that people will perceive causal connections between unrelated events to be of a lower quality (i.e., have poorer concept-coherence) than more straightforward causal scenarios (as in Connell and Keane's study). If this were the case, then ratings in the causal and temporal conditions would be expected to be capped at a level of moderate plausibility.

Method

Materials & Design. The materials and design were the same as in Experiment 1.

Participants. Thirty-six student volunteers from University College Dublin, who had not taken part in Experiment 1, participated in this experiment.

Procedure. The procedure was the same as in Experiment 1, except that participants were asked for a plausibility rating rather than a forced choice judgement. In all conditions, participants were asked "How plausible do you find this scenario?" and asked to circle a rating on a scale from 0 – 10. A rating of 0 was described as meaning the scenario was "not at all plausible", while 5 meant "moderately plausible" and 10 meant "completely plausible".

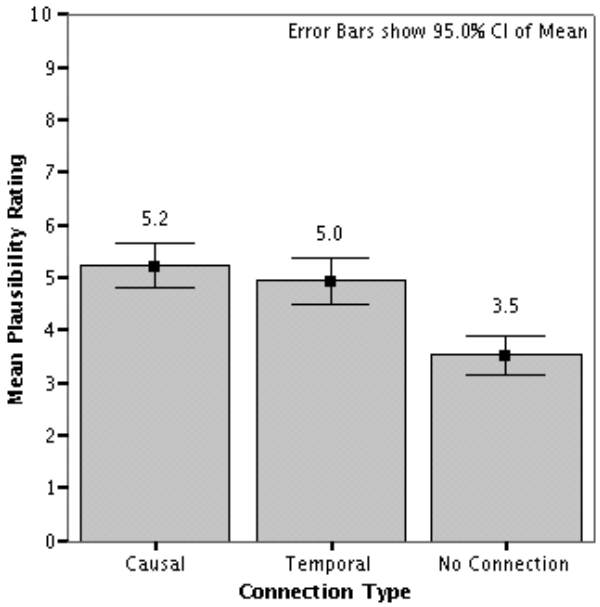

Figure 2: Mean scenario plausibility ratings for each connection type in Experiment 2.

Results & Discussion

The results were in line with predictions, and are shown in Figure 2. As before, three scenarios that were considered plausible by more than 80% of participants in Experiment 1's no-connection control condition were excluded from the analysis in all conditions. Analyses of variance are performed by-participants (F_1) and by-items (F_2), treating participants and items as random factors, respectively.

People's plausibility ratings were influenced by how they had been encouraged to represent the connections between events, as shown by the main effect of connection type, $F_1(2, 33) = 6.30$, $p < 0.005$; $F_2(2, 32) = 31.36$, $p < 0.0001$. Planned comparisons showed that, when asked to represent the events as a specific temporal sequence, people judged the scenarios to be significantly more plausible (5.0, compared to 3.5 control), $F_1(1, 22) = 9.53$, $p < 0.005$; $F_2(1, 16) = 50.58$, $p < 0.0001$. Similarly, when people were asked to represent a causal chain between the events, they perceived the scenarios as being significantly more plausible (5.2, compared to 3.5 control), $F_1(1, 22) = 11.74$, $p < 0.005$; $F_2(1, 16) = 53.23$, $p < 0.0001$. However, there was no difference between the temporal and causal conditions; people did not consider causal

connections between events to be any more plausible than temporal connections, $F_1 < 1$; $F_2(1, 16) = 1.30$, $p > 0.25$.

This experiment shows that, although implausible scenarios may become acceptably plausible in the causal condition, they can never seem *highly* plausible. In other words, while unrelated events can be causally connected, they do not fit with prior knowledge quite as well as more obvious causal connections. For example, we may causally connect the events in the scenario "*The teacher misspelled a word. The vase smashed.*" by assuming that the vase smashed because the teacher bumped against it when taking a step back to examine the misspelled word. While this scenario may just about seem acceptably plausible, it does not seem highly plausible. It certainly does not seem as plausible as a more straightforward causal scenario like "*The cat knocked over a vase. The vase smashed.*" Similarly, for the temporal condition, the results show that connecting events in a specific time frame (e.g., the vase smashed seconds after the teacher misspelled a word) makes a scenario seem somewhat plausible. However, temporal scenarios are considered only moderately plausible at best (Connell & Keane, in press), which suggests that a temporal connection between unrelated events fits with prior knowledge about as well as any other temporal connection. In short, the concept-coherence and plausibility of a normally implausible scenario can be manipulated by encouraging people to make particular connections between events, but the scenario will generally not be judged more than moderately plausible.

So what is the relationship between judging *whether* a scenario is plausible and judging *how* plausible it is? In Experiment 1, we saw that 65% of scenarios in the causal condition were considered acceptably plausible, but yet in Experiment 2, these same scenarios received a plausibility rating of only 5.2 / 10. Similarly, 42% of scenarios in the temporal condition were considered acceptably plausible, and yet were also rated at 5.0 / 10. Analysis of the percentage of plausible responses in Experiments 1 and the mean plausibility ratings in Experiment 2, for each scenario in each condition, shows a direct linear relationship between a scenario's plausibility rating and its acceptability (see Figure 3). This relationship has a significant correlation of $r = 0.88$, $N = 60$, $p < 0.0001$. In short, scenarios with a high plausibility rating will be accepted by most people, while scenarios with a low plausibility rating will be rejected by most people. This suggests that there is no absolute plausibility threshold, above which a scenario will be accepted by everyone as completely plausible. Rather, it depends on what level of acceptability is desired. For example, if we wish 90% of people to accept a scenario, then it should have a mean plausibility rating of approximately 7 out of 10.

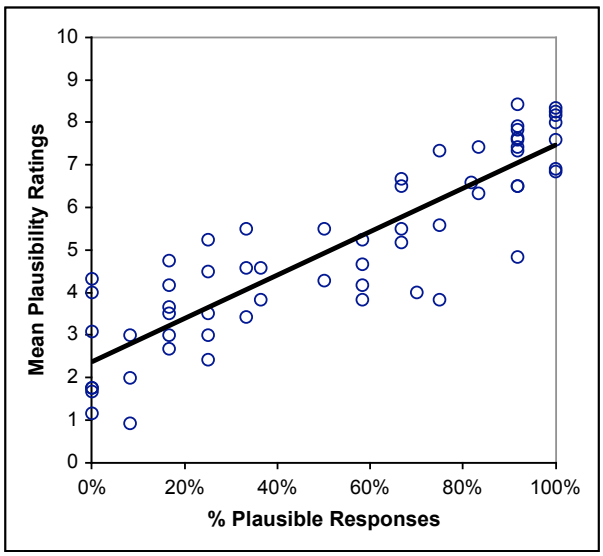

Figure 3: Relationship between scenarios' plausibility ratings and their frequency of acceptance.

General Discussion

This paper shows that an implausible scenario can be made seem plausible by forcing a particular connection between its events. Encouraging people to represent events in a causal chain or temporal sequence, without specifying what the connection should be, alters their perceptions of a scenario's plausibility. In Experiment 1, people are shown to accept 65% of scenarios as plausible once they had explicitly noted a possible causal chain, and 42% of scenarios once they had explicitly noted a possible temporal frame for the events. This compares to a low rate of acceptance for the same scenarios when people are free to make any connections they choose. In Experiment 2, people are shown to consider scenarios as moderately plausible when they are guided into connecting events causally or temporally. In contrast, the same scenarios receive low plausibility ratings when people are free to make any connections they choose. Thus, the novel empirical work reported here demonstrates how people's perceptions of plausibility can be influenced by the circumstances surrounding the task. These findings have implications for any research making use of plausibility judgements, in fields including memory, discourse comprehension, reasoning and conceptual combination.

According to the Knowledge-Fitting Theory of Plausibility (Connell & Keane, 2003a, 2003b, in prep; Connell, 2004), plausibility judgement is about assessing concept-coherence. This view holds that when people make a plausibility judgement, they relate the current scenario to their prior experience, and in some way assess whether it fits in with what they have experienced in the past. Depending on how we

represent a scenario in the first place (i.e., how we connect its events) will therefore determine how well it fits with our prior knowledge. What this paper demonstrates is that the representation of a scenario varies according to how we are encouraged to connect the events. In other words, the concept-coherence and plausibility of a scenario can be manipulated by guiding people towards certain kinds of representation.

The results reported here suggest that people judge plausibility with a certain "cognitive laziness". This means that they do not put any more effort into representing the scenarios than is absolutely necessary. When presented with a scenario, if a possible connection between events does not immediately leap out, then people do not take the trouble to connect the events and instead dismiss the scenario as implausible. However, if circumstances require, people are perfectly capable of connecting even the most disparate events in a coherent manner. For example, the scenario "*The teacher misspelled a word. The vase smashed.*" contains events that are quite difficult to connect. However, people were well able to connect these unrelated events in the causal condition of both experiments, as evinced by the wide and creative variety of causal chains given – e.g., the vase smashed because the teacher bumped against it when stepping back from the blackboard, or because the teacher smashed it in a temper after realising the mistake, or because it was knocked over by a student eager to correct the teacher's error. It is only when circumstances demand it that people overcome their cognitive laziness and take the trouble to reason out a possible connection between events. Indeed, this "cognitive laziness" view is not without its advantages. As well as allowing people to judge plausibility with the least amount of computational expense, it is also tends towards false negatives rather than false positives. This makes it quite a sound approach, as it is safer to reject a scenario that later proves viable than to accept one that later proves unviable.

Acknowledgments

This work was supported by funding from the Irish Research Council for Science, Engineering and Technology, under the Embark Initiative. Thanks also to Dermot Lynott for valuable comments, and to Amy Bohan and Fergus Toolan for experimental assistance.

References

Black, A., Freeman, P., & Johnson-Laird, P. N. (1986). Plausibility and the comprehension of text. *British Journal of Psychology, 77*, 51-60.

Collins, A., & Michalski, R. (1989). The logic of plausible reasoning: A core theory. *Cognitive Science, 13*, 1-49.

Connell, L. & Keane, M. T. (2003a). PAM: A cognitive model of plausibility. *Proceedings of the Twenty-Fifth Annual Conference of the Cognitive Science Society.* Hillsdale, NJ: Erlbaum.

Connell, L., & Keane, M. T. (2003b). The knowledge-fitting theory of plausibility. *Proceedings of the Fourteenth Irish Conference on Artificial Intelligence and Cognitive Science* (pp. 40-45). Ireland: Trinity College Dublin.

Connell, L. (2004). *A cognitive theory and model of plausibility.* Ph.D. thesis, Department of Computer Science, University College Dublin, Ireland.

Connell, L., & Keane, M. T. (in press). What Plausibly Affects Plausibility? Concept-Coherence & Distributional Word-Coherence As Factors Influencing Plausibility Judgements. To appear in *Memory and Cognition*.

Connell, L., & Keane, M. T. (in prep.). The knowledge-fitting theory of plausibility. *Manuscript in preparation*.

Costello, F., & Keane, M. T. (2000). Efficient Creativity: Constraints on conceptual combination. *Cognitive Science, 24*, 299-349.

Keane, M. T. (1997). What makes an analogy difficult?: The effects of order and causal structure in analogical mapping. *Journal of Experimental Psychology: Learning, Memory and Cognition, 23*, 946-967.

Keane, M. T., Connell, L., & O'Donoghue, N. (in prep.). Questioning plausibility judgements: how questions and inferences affect plausibility. *Manuscript in preparation*.

Lemaire, P. & Fayol, M. (1995). When plausibility judgments supersede fact retrieval: The example of the odd-even rule effect in simple arithmetic. *Memory and Cognition, 23*, 34-48.

Lynott, D., Tagalakis, G., & Keane, M. T. (in press). Conceptual combination with PUNC. To appear in *AI Review*.

McKoon, G. & Ratcliff, R. (1992). Inference during reading. *Psychological Review, 99*, 440-466.

Pickering, M. J., & Traxler, M. J. (1998). Plausibility and recovery from garden paths: An eye-tracking study. *Journal of Experimental Psychology: Learning, Memory and Cognition, 24*, 940-961.

Reder, L. M. (1982). Plausibility judgments vs. fact retrieval: Alternative strategies for sentence verification. *Psychological Review, 89*(3), 250-280.

Reder, L. M., Wible, C., & Martin, J. (1986). Differential memory changes with age: Exact retrieval versus plausible inference. *Journal of Experimental Psychology: Learning, Memory & Cognition, 12*(1), 72-81.

Smith, E. E., Shafir, E., & Osherson, D. (1993). Similarity, plausibility, and judgments of probability. *Cognition, 49*, 67-96.

Speer, S. R., & Clifton, C. (1998). Plausibility and argument structure in sentence comprehension. *Memory and Cognition, 26*, 965-978.