Expecting a Surprise?
The Effect of Expectations on Perceived Surprise in Stories

Rebecca Grimes-Maguire (rebecca.grimes@ucd.ie)
Mark T. Keane (mark.keane@ucd.ie)
Department of Computer Science, University College Dublin
Belfield, Dublin 4, Ireland

Abstract

This paper describes two novel experiments that investigate the cognitive basis of surprise in stories depicting simple, everyday events. Participants were asked to read a number of short scenarios, each of which concluded with some surprise event for the central protagonist. The story versions differed in how strongly they predicted this 'surprise' conclusion, by varying the degree of attention drawn to key enabling conditions for that conclusion. The effect of this manipulation on participants' surprise ratings, and the speed with which they read the final sentence, were used as dependent measures. The results of our experiments show that the specific representations built by people in understanding the earlier part of a story have definite effects on their level of surprise at later events. Furthermore, the pattern of reading times for the target sentences supports the explanation given for these differences. We discuss the consistency of these findings with theories of discourse comprehension and describe how such effects might be modelled computationally.

Keywords: Discourse Comprehension; Surprise; Expectation.

Introduction

As well as being one of the most basic and universal of emotions, surprise has long been recognised as a fundamental building block of human cognition (Darwin, 1872; Ekman, 1972). Despite this, cognitive investigations of surprise, which move beyond the physiological and affective domain, have only recently emerged (e.g. Fisk, 2002; Schützwohl, 1998; Teigen & Keren, 2003). The occurrence of surprise in narrative comprehension presents a good starting point for tackling this complex phenomenon, as it is undoubtedly an important and central aspect of many stories. Surprise, for instance, has been shown to be vital in determining the perceived interestingness of a piece of text (e.g. Iran-Nejad, 1987; Kim, 1999); very few best-sellers describe the hero going to the office, working all day and then coming home and going to bed. Instead, writers often maintain a reader's interest with well-timed surprises; for instance, Raymond Chandler's murder discovery scenes often describe the dead body obliquely (e.g., a brown loafer protruding from the curtain), so that the reader's surprise at discovering the referent (i.e., that the shoe is attached to a corpse) mirrors the detective's surprise at discovering the body.

In this paper, we try to elucidate the cognitive basis of surprise in understanding narrative. Evidently, a person should be surprised when something unexpected occurs in an event sequence. However, as discussed in the next section, the relationship between expectation and surprise is not quite so clear-cut. We examine the hypothesis that a surprising event is unexpected relative to a reader's mental representation of the sequence. In other words, when reading a story, a person's level of surprise should be determined by whether or not the specific conditions *supporting* the surprising event were built into their representation as the prior narrative was comprehended.

Surprise and Expectations

At a cognitive level, surprise is intuitively related to expectations, and as such the Oxford English dictionary defines it as *"an unexpected occurrence or event"*. However, this may not be a straightforward relationship due to differing interpretations of what it means to expect something. Teigen & Keren (2003) observe that in cognitive psychology, expectations are usually quantified in terms of probabilities; i.e. an unexpected event is considered to be a low probability outcome. Yet, events that have an equal likelihood of occurring are not always judged to be equally surprising by an observer (e.g. Schackle, 1969). Surprise then, might not be due to the low probability of an outcome, but to its degree of contrast with a more likely, but disconfirmed, expectation. The Expectancy-Disconfirmation hypothesis, (Meyer Reisenzein & Schützwohl, 1997), similarly suggests that surprise is experienced when expectations based on preexisting schemas fail to be confirmed. If we apply these ideas to the area of discourse comprehension, then surprise is likely to occur when a predicted sequence of events in a story is violated or deviated from in a significant way.

While little research has investigated the conditions which evoke surprise in narrative, a lot of work has examined people's ability to form *forward inferences* during reading (e.g., Klin et al., 1999; McKoon & Ratcliff, 1986). These inferences operate by linking representations of the text with background knowledge to predict a likely outcome, thus facilitating the comprehension process. It seems feasible that when people read a familiar event sequence in a story, they are able to make predictions about what will happen next. According to Schank & Abelson (1977), this is done by means of script-based schemas. For example, when reading the scenario *"The cup of coffee was balanced on the arm of the chair. Suddenly, Richard sneezed."*, most people can predict that Richard's sneeze will cause the cup of coffee to spill or be knocked off the chair. If the reader subsequently discovers that a cat on Richard's knee jumps up and scratches him, they will probably be surprised at this ending, as, despite

being plausible, it does not fit in well with the expected causal sequence.

Grimes-Maguire & Keane (2004) asked people to provide either a predictable or a creative ending to simple two-sentence scenarios like the above example. It was found that even when participants were told to be creative, they were constrained by their background knowledge, as the endings they gave to the scenarios overlapped significantly with those generated by people asked to give predictable endings. This suggests that participants found it quite difficult to avoid considering the typical outcomes to such sequences, even when they were instructed to provide novel endings.

However, Grimes-Maguire & Keane (2004) explicitly asked participants to form an expectation, so this work does not address the issue of whether readers *automatically* anticipate events during reading, i.e. by making a forward inference. Some researchers propose that such 'on-line' expectations are rarely formed, because making predictions about events that may never happen is an uneconomical reading strategy (McKoon & Ratcliff, 1986). Nevertheless, a wealth of empirical evidence has suggested that under certain conditions people do make forward inferences (e.g. Calvo, 2000; Fincher-Kiefer, 1996; Klin et al., 1999). The emerging consensus today is that readers can anticipate outcomes, but only in highly predictable scenarios. For example, when reading the sentence *"He threw the delicate porcelain vase at the wall"* most will infer that *"the vase broke"* (Klin et al, 1999).

If we assume a direct relationship between expectation and surprise, then readers should only experience surprise when an explicitly represented, highly predictable event in a discourse fails to subsequently occur (leading to a disconfirmed expectation). This also means that the degree to which a reader's story representation supports a prediction, as mediated by events in the preceding discourse, should critically influence their level of perceived surprise at future unfolding events. This is the hypothesis tested by the current research.

Expecting a Surprise

In examining this hypothesis, there is one further issue that is unique to surprise in discourse comprehension. Readers who form expectations in the story context are probably operating off at least two distinct types of knowledge:

(i) Knowledge about the causal structure of the world derived from direct and vicarious experience (e.g. Trabasso & van den Broek, 1985)

(ii) Knowledge of story structure (e.g. Stein & Glenn, 1979).

It seems likely that these two types of knowledge drive comprehension, influencing the reader's representation, and determining whether or not a forward inference will be made. Rapp & Gerrig (2002) have provided evidence for a comparable distinction between 'reality-driven' and 'plot-driven' analysis in reading. In practice, this means that in the story context, readers are more sensitive to depicted details that may set up subsequent surprises. The author has the power to control the amount of information given to the reader about an upcoming surprise, and subtle differences in narration can have a great impact on the reader's representation of events. For example in murder mysteries, if a room is described as having a gun hanging over the mantelpiece, there is often an inevitability about its subsequent use. We could say that in such cases, readers can 'predict' the surprise, since the enabling conditions for an expectation have been built into their representation.

Overview of Experiments

We wished to establish whether a reader's surprise is related to the degree to which they can represent expectations about the end event in a story. Few previous studies have examined both of these factors together. We devised a variety of short scenarios which differed in how strongly they set up the final, 'surprising' sentence (Experiment 1). The dependent variables were the time taken to read the concluding sentence (to determine whether the appropriate forward inference had been made) and a rating of surprise for the ending. Three variants of the story scenario were used: Predictable, Neutral and Unpredictable. We hypothesised that participants would register progressively more surprise across these conditions in the order shown, even though we envisaged that the relevant explicit forward inference would only be made in the Predictable condition. We also examined what happened when an enabling condition for the concluding sentence was 'undone' (Exp. 2)

Experiment 1

In this experiment we manipulated the amount of information that was given to the reader about an upcoming surprise for the central protagonist. We did this by creating short scenarios in which only the second sentence was changed in each version. This sentence varied in how strongly it predicted the story ending. There were three different conditions: 1) *Predictable*, in which the sentence described a specific event for enabling the surprise conclusion for the protagonist, 2) *Neutral*, in which the sentence hinted at the upcoming surprise by containing vague information about the enabling event, and 3) *Unpredictable*, which did not include any information about the enabling condition, instead describing some irrelevant event. The difference between these three conditions is illustrated in the sample scenario in Table 1.

Pre-test

Before the main experiment, it was necessary to carry out a pre-test in order to ensure that the three conditions differed significantly in how strongly they predicted the story conclusion. Fifteen participants were presented with a booklet containing 20 short stories in each of their three versions (Predictable, Neutral and Unpredictable). Participants were then instructed to rate each version of the story for how strongly they felt it predicted the final sentence on a scale of 1-7 (with 1 denoting a very unpredictable ending and 7 denoting a very predictable ending). A one-way, repeated

measures ANOVA revealed a significant effect of condition in these ratings both by-participants ($F_1(2,28) = 241.069$; $p<0.001$; $MSe = 0.291$), and by-materials ($F_2(2,38) = 378.230$; $p<0.001$; $MSe = 0.247$). As expected, the Predictable condition received the highest ratings ($M = 6.05$) followed by the Neutral ($M = 4.22$) and the Unpredictable ($M = 1.74$) conditions.

Method

Participants Thirty-six people (aged between 20-58 yrs) voluntarily participated in the experiment.

Materials Eighteen short stories were randomly selected from the original set of 20 used in the pre-test. Each story was five sentences long with the fifth sentence depicting an ending to the scenario. The different conditions were determined by changing the second sentence accordingly, as can be seen in Table 1. The final sentence in each story was modified to be 10-11 syllables long, thus allowing us to make reliable comparisons of reading times across the materials.

Table 1: Sample scenario from Experiment 1

John was drinking coffee in the sitting room.

The cup was balanced on the armchair (Predictable).
He put the cup of coffee down (Neutral).
He started to read the paper (Unpredictable).

He wasn't feeling very well.
Suddenly he sneezed.
The cup of coffee spilt all over the carpet.

Design A Within-Participants design was used for the main experimental manipulation of condition (or story version). Each participant read a randomly selected set of eighteen stories: six Predictable, six Neutral and six Unpredictable.

Procedure and Scoring Instructions for the task were displayed on a computer screen, as follows: "You will be asked to read a number of short stories (each with 5 sentences). These will be presented sentence-by-sentence on the screen. When you have finished reading a sentence, press the spacebar to display the next sentence. Please read each story carefully. At the end of each story you will be asked how surprising you found the final sentence. You can do this by choosing a number between 1 and 7, where 1 means a very unsurprising sentence and 7 means a highly surprising sentence." Below these instructions was a depiction of a 7-point scale, which ranged from 1 (*Very Unsurprising*) to 7 (*Highly Surprising*). After reading the instructions, the participants were given a practice trial, which consisted of two stories in the same format as the main experiment. On completion of the practice trial, they commenced the experiment. The stories were presented in a different random order to each participant.

Error responses were excluded from the data prior to

analysis (e.g., miss-keyed presses). Also, in the reading time data, responses deemed too fast (i.e. below 300ms) or too slow (i.e. above 10000ms) were omitted. Any responses that were above 3 standard deviations of a participant's mean response time were excluded as well. This resulted in a total of 5% of the data being omitted from the analysis.

Results

As expected, participants' surprise ratings systematically varied with the degree of attention given to the enabling condition for the surprising event. Furthermore, the reading time data supports the proposal that forward inferences are only made in predictable situations.

Surprise ratings Figure 1 shows the mean surprise ratings given across the three conditions. A one-way, repeated measures ANOVA revealed a significant difference between these conditions, using both participants ($F_1(2,70) = 60.518$; $p < 0.001$; $MSe = 0.407$) and materials ($F_2(2,34) = 21.353$; $p < 0.001$; $MSe = 0.563$) as random factors. Post-hoc analysis showed that all three conditions differed significantly; the endings of the Predictable stories were judged as the least surprising ($M = 2.806$), followed by the Neutral ($M = 3.819$) and then the Unpredictable ($M = 4.444$) story endings. Non-parametric analysis using Page's trend test revealed a reliable trend for the Predictable < Neutral < Unpredictable ordering of conditions by-participants ($L_1(2) = 487$; $p <0.001$), and by-materials ($L_2(2) = 238$; $p <0.001$).
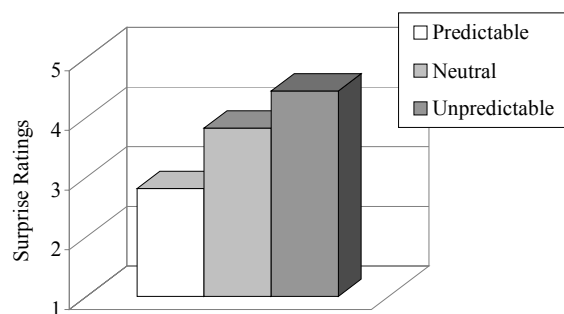

Figure 1: Mean surprise ratings given across Predictable, Neutral and Unpredictable story versions (Exp.1)

Reading times for concluding sentence It was expected that the reading times for the story ending would reflect forward inferences made about the surprising event. If the early sentences strongly suggest a particular outcome, then that outcome should be partially or wholly represented, making reading times for the concluding sentence faster. A one-way ANOVA repeated measures did reveal a reliable main effect of condition by-participants ($F_1(2,70) = 4.971$; $p<0.05$; $MSe = 321627.5$), but not by-materials ($F_2(2,34) = 2.582$; $p>0.05$; $MSe = 291331.8$). However, pairwise comparisons using a one-tailed dependent t-test revealed a significant difference between the Predictable and Neutral conditions across both participants ($t_1(35) = 2.5216$; $p <0.01$) and materials ($t_2(17) = 2.076$; $p <0.05$). This difference was also observed between the Predictable and Unpredictable conditions ($t_1(35) = 3.189$;

$p<0.01$; $t_2(17) = 1.799$; $p<0.05$). The Neutral and Unpredictable conditions did not differ. This shows that overall, the mean reading times for the Predictable condition ($M = 3106.05$) were reliably different to both the Neutral ($M = 3298.31$) and Unpredictable ($M = 3445.31$) conditions.
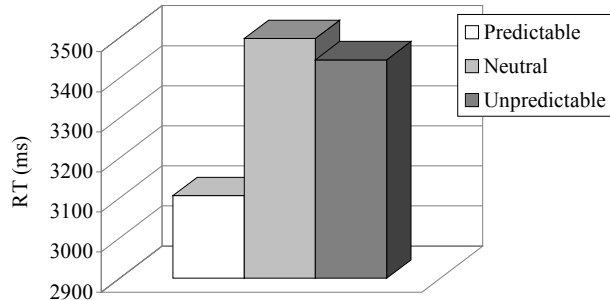

Figure 2: Reading times for final sentence across the Predictable, Neutral and Unpredictable conditions (Exp.1)

Discussion

The results of Experiment 1 show that varying the type of information given about an enabling event for the story conclusion had a significant impact on the reader's representation of the scenario. The three different conditions resulted in qualitatively different surprise ratings for the same ending. Also readers were faster to read the conclusion in the Predictable condition, which suggests that they were making forward inferences about what was about to happen. This result is notable given that the only difference between the stories was the second sentence, which had the same number of syllables in each condition and which contained no contradictory information across the three versions.

The most interesting finding of this experiment is that surprise and expectancy do not appear to have a linear relationship. While readers only seem to form an 'on-line' expectation in highly predictable scenarios (as illustrated by faster reading times of the conclusion), they can nevertheless experience qualitatively different levels of surprise across the three conditions.

Experiment 2

In Experiment 1, it is worth noting that the surprise ratings given for the Neutral condition (which provides a generalised enabling condition for the surprise event) sit roughly midway between those given for the Predictable condition (which provides the key enabling condition) and the Unpredictable condition (which gives little indication of an enabling condition). Presumably, reading stories in the Neutral version leads to some representation of the enabling condition for the outcome event, hence the reduced level of surprise. If this proposal is true, then it should be possible to increase perceived surprise in the Neutral condition if we explicitly *undo* that enabling event. For example, by saying *"John went to turn on the TV"* we might imply to the reader that John has

moved away from the cup of coffee in order to turn on the television (see Table 2). In our second experiment, we tested this hypothesis by examining the effects of such an 'Undoing' condition with the Predictable and Neutral conditions. We expected that the increasing trend in surprise should be as follows: Predictable < Neutral < Undoing. On the issue of reading times for the concluding sentence, as in Experiment 1, we expected the Predictable condition to yield faster times than both the Neutral and Undoing conditions, because the representation of stories in this condition should induce forward inferences.

Table 2: Sample scenario for Experiment 2

John was drinking coffee in the sitting room.

The cup was balanced on the armchair.
He started to read the paper. (Predictable).
He put the cup of coffee down.
He started to read the paper. (Neutral).
He put the cup of coffee down.
He went to turn on the TV. (Undoing).

He wasn't feeling very well.

Suddenly, he sneezed.

** The cup of coffee spilt all over the carpet. **

Comprehension Question: Was John feeling sick?

Method

Participants Fifty-nine undergraduate Computer Science students from University College Dublin (17-20 yrs) participated in this experiment for partial course credit.

Materials Fifteen stories were used based on the stories in Experiment 1 (the remaining three were difficult to modify to meet the constraints of this experiment). Each story had six sentences, where the second and third sentence varied in each condition. In the Predictable and Neutral conditions the second sentence was the same as the analogous conditions in the previous experiment, while the third sentence included irrelevant information that did not relate to the story conclusion (this was to ensure that all versions of the story were the same length). In the Undoing condition, sentence 2 was the same as in the Neutral version of the story, while sentence 3 undid the events of the second sentence, as can be seen in Table 2. We also included a comprehension question for each scenario to ensure participants would read the story carefully.

Design As in Experiment 1, the design was Within-Participants, where each participant was randomly assigned to read fifteen stories: five Predictable, five Neutral and five Undoing scenarios.

Procedure and Scoring The procedure was the same as that used in Experiment 1, with the addition that participants were

also asked a simple question about each story after making their surprise rating. This question required a yes or a no answer, indicated by pressing Y or N on the keyboard. As before, participants were given a practice session and all materials were randomly ordered. Following the procedure in the first experiment, 3.1% of the total data was discarded prior to analysis.

Results and Discussion

The predictions made for this experiment were confirmed by the expected trend in surprise ratings and the pattern of reading times for the final sentence (see Figures 3 and 4).

Surprise ratings A one-way, repeated measures ANOVA examining the surprise ratings showed a main effect of condition, using both participants ($F_1(2,116)$ = 47.7608; MSe = 0.738; $p<0.0001$), and materials ($F_2(2,28)$ = 13.989; MSe = 0.627; $p<0.0001$) as random factors. Page's trend test also revealed a significant trend for the ordering across the conditions (Predictable < Neutral <Undoing) for participants ($L_1(2)$ = 772.5; $p<0.001$), and materials ($L_2(2)$ = 199; p <0.001). As in Experiment 1, the Predictable condition induced the lowest surprise rating (M = 3.07) followed by the Neutral (M= 4.27) and Undoing (M = 4.51) conditions. Overall, it should be noted that while the Undoing condition conforms to the predicted trend, it is not as markedly different to the Neutral condition as the Predictable one.
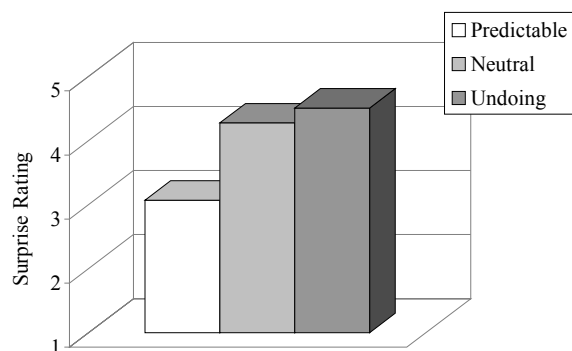

Figure 3: Mean surprise ratings given across Predictable, Neutral and Undoing conditions (Exp. 2)

Reading times for conclusion A one-way, repeated measures ANOVA revealed a main effect of condition for reading times of the concluding sentence, by-participants ($F_1(2,116)$ = 7.456; MSe = 158614.6; $p<0.001$), and by-materials ($F_2(2,28)$ = 4.239; MSe = 72031.61; p <0.05). Post-hoc analysis showed that the Predictable (M = 1867.08) differed significantly from the Neutral (M= 2102.654) and the Undoing (M = 2125.468) conditions, but the latter two conditions were not reliably different. This result follows that of Experiment 1, suggesting that participants only make inferences in predictable situations.
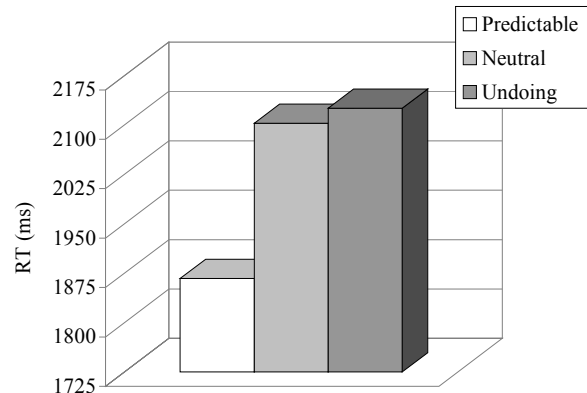

Figure 4: Reading time for final sentence for the Predictable, Neutral and Undoing conditions (Exp. 2)

General Discussion

The aim of this research was to explore the cognitive basis of surprise in understanding simple stories. The results indicate that target events will be judged as surprising relative to the degree to which they are specifically supported in a reader's representation. This is achieved by the presence of enabling conditions and forward inferences. In our Predictable story versions, an explicit forward inference appears to be made about the concluding event, making it less surprising, and quicker to read. In our Neutral versions, a forward inference does not appear to be made (reading times are much slower), but the presence of implicit enabling conditions in the representation impacts the surprise ratings given. When these enabling conditions are absent (as in the Unpredictable condition), or undone (as in the Undoing condition), overall surprise ratings are the highest. These results are quite consistent with models of comprehension and are currently the focus of computational modelling work being carried out by the authors.

A Changing Representation

The findings of these experiments illustrate that the type of information presented in a story can lead the reader to build a unique representation of the depicted events. This result is supportive of the Situation Model of comprehension (Zwaan & Radvansky, 1998), which holds that when presented with a story, readers do not just passively absorb what is on the page in front of them, but instead build up a rich representation of the depicted scenario. They make assumptions about the central characters, their goals and actions, and even form a mental picture of the location and time in which the story is set. This detailed representation is achieved by calling upon the wealth of background knowledge readers have about the world, and is highly sensitive to the way in which the information is presented. We have also seen that, in turn, this representation strongly determines whether subsequent events will appear to be surprising or not, based on the degree to which they are supported by prior events in the discourse.

Our results illustrate that while anything that is expected will be judged as unsurprising, not everything unexpected will be judged as equally surprising. We found that even though explicit expectations were only formed in the Predictable condition, readers in both experiments could in fact distinguish between the various levels of predictability depicted across the three story versions (as was evident in the different surprise ratings). This suggests that expectation and surprise are represented differently in the reader's mind.

We have begun to model these effects in an extension of a previous model of plausibility judgements (see Connell & Keane, 2003; 2004), which creates inference paths in order to connect the events in a discourse. The key question for our extension of this model is how to relate the representations that people build to the surprise ratings that they generate for certain events. Based on the results of our experiments, it would seem that there are two separate processes taking place: One process occurs *during* reading and involves building an up-to-date representation of the narrative situation. This process controls the formation of forward inferences when the discourse is highly suggestive of a particular outcome. The other process occurs *after* reading, when the reader must make a surprise judgement. Here, they must assess their representation in order to search for any information (i.e. enabling conditions) that is consistent with the surprising event. By formalising these two processes, we hope to make clear the distinction between expectation formation and surprise in narrative.

Concluding Comments

In this paper we have outlined two novel experiments that examine the interplay between predictability and surprise in reading. This research highlights the fact that text comprehension is dependent not only on our background knowledge of the real world, but on the expectations we have regarding story structure. People often 'expect' surprises when reading, but the degree to which they do this is mediated by the amount of information provided to them by the author. By better understanding the nature of this, we hope to achieve a more in-depth understanding about the complex nature of event representation in discourse.

References

Calvo, M.G. (2000). The time course of predictive inferences depends on contextual constraints. *Language and Cognitive Processes, 15,* 293–319.

Connell, L. & Keane, M.T. (2003). 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. (2004). What plausibly affects plausibility: Concept coherence and distributional word coherence as factors influencing plausibility judgements. *Memory & Cognition, 32 (2),* 185–197.

Darwin, C. (1872). *The Expression of Emotion in Man and Animals.* London: J. Murray.

Ekman, P. (1972). *Emotion in the Human Face: Guidelines for research and an integration of findings.* New York: Lawrence Erlbaum.

Fincher-Kiefer, R. (1996). Encoding differences between bridging and predictive inferences. *Discourse Processes, 22,* 225–246.

Fisk, J.E. (2002). Judgments under uncertainty: Representativeness or potential surprise? *British Journal of Psychology, 93,* 431-449.

Grimes-Maguire, R. & Keane, M.T. (2004). Evidence of Muddy Knowledge in Reaching for the Stars: Creating Novel Endings to Event Sequences. *Proceedings of the Twenty-Sixth Annual Conference of the Cognitive Science Society.* Hillsdale, NJ: Erlbaum.

Iran-Nejad, A. (1987). Cognitive and affective causes of interest and liking. *Journal of Educational Psychology*, 79, 120-130.

Kim, S-I. (1999). Causal bridging inference: A cause of story interestingness. *British Journal of Psychology, 90*, 57-71.

Klin, C.M., Murray, J.D., Levine, W.H., & Guzmán, A.E. (1999). Forward inferences: From activation to long-term memory. *Discourse Processes, 27*, 241-260.

McKoon, G. & Ratcliff, R. (1986). Inferences about predictable events. *Journal of Experimental Psychology. Learning, Memory, and Cognition, 12,* 82–91.

Meyer, W.U., Reisenzein, R. & Schützwohl, A. (1997). Towards a process analysis of emotions: The case of surprise. *Motivation and Emotion, 21(3),* 251-274.

Rapp, D.N., & Gerrig, R.J. (2002). Readers' reality-driven and plot-driven analyses in narrative comprehension. *Memory & Cognition, 30*, 779-788.

Schank, R.C. & Abelson, R.P. (1977). *Scripts, Plans, Goals, and Understanding: An Inquiry into Human Knowledge Structures.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Shackle, G.L.S (1969). *Decisions, Order and Time in Human Affairs*. Cambridge: Cambridge University Press.

Stein, N.L. & Glenn, C.G. (1979). An analysis of story comprehension in elementary school children. In: R.O. Freedle (ed.), *New Directions in Discourse Processing, vol 2.* Norwood, NJ: Ablex Publishing Corporation.

Schützwohl, A. (1998). Surprise and schema strength. *Journal of Experimental Psychology, Learning, Memory and Cognition, 24,* 1182-1199.

Teigen, K.H., & Keren, G. (2003). Surprises: Low probabilities or high contrasts? *Cognition, 87*, 55-71.

Trabasso, T. & van den Broek, P.W. (1985). Causal thinking and the representation of narrative events. *Journal of Memory and Language, 24,* 612-630.

Zwaan, R.A. & Radvansky, G.A. (1998). Situation Models in Language Comprehension and Memory. *Psychological Bulletin, 123 (2),* 162-185.