A Cognitive Model of Surprise Judgements

Rebecca Maguire (rebecca.grimes@ucd.ie)
Fintan Costello (fintan.costello@ucd.ie)
Mark T. Keane (mark.keane@ucd.ie)
School of Computer Science and Informatics, University College Dublin
Belfield, Dublin 4, Ireland.

Abstract

In this paper we outline a cognitive theory and model of surprise judgements which aims to explain how and why some events are considered to be surprising in a piece of text, while others are not. The model is based on a series of experiments carried out by Grimes-Maguire and Keane (2005a), which show that subtle changes in the predictability of a discourse can have a profound effect on a reader's perceived surprise at certain events. Rather than defining surprise in terms of expectation, we conceive of it as a process involving *Representation-Fit*. We have implemented this theory in a computational model that has two stages: the Integration stage entails building a coherent representation of the scenario by means of an objective knowledge base rooted in WordNet. The Analysis stage then outputs a surprise rating for a specified event, based on the degree to which that event can be supported by the prior representation. Simulations reveal a strong correspondence between model and participant generated surprise ratings.

Keywords: Surprise, cognitive modelling, representation.

Introduction

Although we have a remarkable ability to make sense of and even predict events in the world around us, sometimes this ability breaks down and we experience a feeling of surprise. Consider how you would react, for instance, if you heard a sudden loud bang while sitting in a quiet room, or how you might feel if you saw your next-door neighbour while on holiday in another country. It is well established that the perception of such surprising events will often initiate a number of complex cognitive processes, usually leading to an interruption of ongoing activities and an increased focusing of attention on the event in question (e.g. Ekman, 1972; Fisk, 2002; Meyer, Reisenzein, and Schützwohl, 1997). The main purpose of these actions is to try and understand why the surprising event occurred in the first place, so as to resolve any feelings of confusion.

Despite our familiarity with the concept of surprise however, it is quite difficult to reach a satisfactory definition of this in the literature (see Maguire & Keane, 2006). Intuitively, we can relate surprise to expectation or probability, but complications arise here in that not *every* low probability event will be judged as equally surprising by an observer, and similarly there may be some high probability events that are nevertheless judged as surprising (e.g. Shackle, 1969; Teigen & Keren, 2003). While some research has recently identified a link between the *subjective* probability of a given event occurring and its surprise level

(Fisk, 2002), other theorists claim that this phenomenon may be best defined in terms of expectations, or more specifically, disconfirmed expectations. For example, Meyer et al (1997) in their *Cognitive-Psychoevolutionary Model,* maintain that surprise occurs when an event is seen to deviate significantly from an expected schema. In a similar vein, Teigen and Keren (2003) in their *Contrast Hypothesis* propose that a person's level of surprise at a given event will be dependent on the degree of perceived disparity between that event and another more likely one. In other words, these theories hold that while expectation is vital in the experience of surprise, genuine surprise results only when the event in question conflicts with *another* event that was more expected. The question for us is whether this mechanism is truly the best predictor of surprise, or if there might be something other than expectation at work.

In this paper, we shall address these issues and in doing so will propose a novel cognitive theory and model of surprise judgements. In short, this theory suggests that surprise is not dependent on expectation, but rather on the characteristics of a person's scenario representation. Before detailing the implementation of this model however, we will first outline some empirical work which led to its development.

Predictability, expectation and surprise

While the above research implies that the relationship between expectation and surprise is not quite as clear-cut as initially thought, few studies have examined the connection between these two variables in discourse comprehension. For this reason, Grimes-Maguire and Keane (2005a) carried out a number of experiments to ascertain whether perceived surprise at a given event in a discourse would be related to the degree to which that event was expected, as measured by forward inferences. Forward inferences operate by connecting the events in a text with background knowledge so that the reader can form an expectation about an upcoming event (e.g., if you read the sentence *"He threw the delicate porcelain vase against the wall"*, you will probably infer that *"the vase broke"*). Inferred events will be read quicker, thereby offering a more objective means of quantifying expectation (Klin, Murray, Levine & Guzmán, 1999). In Grimes-Maguire and Keane's (2005a) study, participants were required to read a number of short scenarios, such as that in Table 1, and asked to indicate how surprising they found the concluding sentence. The degree of information given to the participants about this critical sentence was varied in three distinct versions of the

scenarios: in short, the *predictable* version described an explicit enabling event for the final sentence, the *neutral* version merely hinted at this enabling event by containing vaguely supportive information for it, while the *unpredictable* version described an event that was irrelevant to the final sentence.

Table 1: Sample scenario from Grimes-Maguire & Keane, (2005a, Experiment 1)

John was drinking coffee in the sitting room.
The cup 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 coffee spilt all over the carpet.

Table 2: Results from Grimes-Maguire & Keane (2005a)

Condition	Surprise rating		Reading time (ms)	
	Mean	SD	Mean	SD
Predictable	2.81	1.67	3106	1610
Neutral	3.82	1.84	3498	1839
Unpredictable	4.44	1.82	3445	1608

The results of this experiment are displayed in Table 2. We discovered that participants were easily able to distinguish between the separate levels of predictability across the three scenario versions, in so far as they gave qualitatively different surprise ratings for each. However, we also found that participants only formed an on-line expectation in the predictable condition, as indicated by faster reading times of the final sentence relative to the other two conditions. This suggests that expectations could only be made about an upcoming event when the preceding discourse was highly suggestive of it; a finding in line with other studies on forward inferences (e.g. Klin et al, 1999). These results argue against the claim that surprise is directly related to expectation, because if these two concepts were linearly related, participants would have registered low surprise in the predictable condition (when an expectation was formed), but an equally high level of surprise in the other two conditions (where no expectation was formed). While the former was found to be true, the latter was not. To explain these findings, we conceptualise surprise in terms of *Representation-Fit* – a theory outlined in more detail below.

Theory of representation-fit

Most researchers in discourse comprehension agree that, in order to fully understand a piece of text, the reader must construct a number of complex levels of representation (e.g. Graesser, Singer & Trabasso, 1994; van Dijk & Kintsch, 1983). While the more basic of these levels involves grasping the individual meanings of the constituent words and their grouping or syntax, the most sophisticated level entails building a *situation model* of the events in question.

When a reader creates a situation model, it is almost as if they are a direct observer of the depicted events: they can make inferences about the central characters, their goals and actions, as well as forming a mental picture of the time and location in which the story is set (Zwaan & Radvansky, 1998). In the majority of discourses, it is relatively easy for a reader to build such a situation model, mainly because they will be motivated to achieve a sense of *coherence* among the component events by relating every new event to what has gone before. Graesser et al (1994) note that, in doing this, not only is it essential to attain *local coherence* by linking neighbouring events together, (e.g., anaphoric reference - knowing that in the sentences *"John was drinking coffee in the sitting room. He put the cup down"*, *HE* in the second sentence refers to *JOHN* in the first), it is also important to accomplish *global coherence,* by integrating all the events in the text together to make sense of the story. If events can be linked effortlessly and without ambiguity, then the story can be said to be coherent, allowing the reader to build an accurate representation, or situation model, of the events.

In their *Plausibility Analysis Model*, Connell and Keane (2006) established that coherence is vital in determining the plausibility of short event descriptions. Namely, they found that the easier it is to make an inference between two events in a text, the more plausible those events will appear. Given the strong link between surprise and plausibility (Black, Freeman & Johnson-Laird, 1986), it follows that coherence must also play a key role in this phenomenon More specifically, it seems reasonable to suppose that an event which *cannot* be coherently linked with one's scenario representation will be judged as surprising, whereas an event that *can* be coherently linked will result in little or no surprise. This is the central premise of our theory of Representation-Fit (Grimes-Maguire & Keane, 2005b). In short, this theory conceives a judgment of surprise for a given event in a scenario as an attempt to 'fit' that event with the prior discourse representation in the same way as one might attempt to position a piece into a jigsaw puzzle.

The main way in which this account differs from previous theories of surprise is that it does not view expectation as a vital determinant of this experience (e.g. Meyer et al, 1997). Instead, we see the assessment of surprise as consisting of two distinct stages. Firstly, the *Integration stage* involves linking each new event with those that have gone before so as to achieve an up-to-date coherent representation of the scenario. Secondly, the *Analysis stage* involves a systematic assessment of this representation, whereby the reader is required to rate their surprise for a given event. As well as detecting factors that are directly supportive of this event, we propose that readers are also able to identify vaguely supportive information for it. This would explain why surprise ratings differed markedly between the three conditions in the experiment of Grimes-Maguire and Keane (2005a). In the remaining sections we will detail how these distinct stages have been implemented computationally, using the scenario in Table 1 as an example.

Computational Implementation of Model

Based on the theory of representation-fit, we have created a computational model of surprise which takes as input short scenarios such as that in Table 1, and outputs a surprise rating for the final sentence. This is achieved by a number of different components, the most fundamental of which is the knowledge base (KB). While many computational models in discourse comprehension employ hand-crafted knowledge bases, we have chosen to use WordNet as a foundation for ours (cf. Miller, 1995). WordNet is a semantic lexicon for the English language, and can therefore provide a more objective means for representing the information necessary to understand the experimental scenarios. The KB comprises the definitions (or glosses) of the component words[1], the hierarchical relationships between these words, and the acceptable arguments for any verbs or actions which are used in the scenarios. This information was extracted in propositional format from the Prolog implementation of WordNet. The KB also contains some key attributes of the concepts relevant to the scenarios which are not present in WordNet. Such knowledge was incorporated into the KB in a blind fashion and only a small percentage of the propositions are of this nature.
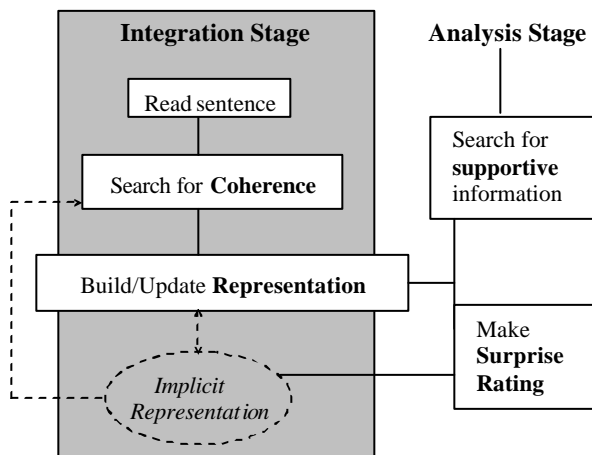

Figure 1: Processes involved in the model of surprise

The two chief components involved in the representation of events are: (1) searching for coherence among the text constituents, and (2) building/updating the representation of the scenario. These are governed by different aspects of the model, and roughly correspond to van Dijk and Kintsch's (1983) distinction between the *textbase* and the *situation model* of comprehension. The component necessary for generating the surprise rating involves detecting the degree of representational support for the critical event by means

of a retrospective judgement. Figure 1 illustrates how these components are implemented in the two different stages.

Integration stage

In the Integration stage, the program takes as input a scenario, sentence-by-sentence (in propositional format), in the same way a reader would. Each sentence must first be deemed coherent, based on background knowledge, before it can be integrated into the representation. For example, the opening sentence in Table 1, *"John was drinking coffee in the sitting room",* is considered coherent because the verb *to drink* is correctly paired with an animate object (*John*) and a liquid (*coffee*). This rule is adapted from the definition of the appropriate sense of 'drink' in WordNet (i.e., *"to take in liquids"*). Since actions like drinking can happen in locations (i.e. *sitting room*), the model can classify this sentence as coherent and proceed to build a representation of the depicted events. This is achieved by using the available dimensions of the *Event-Indexing Model* (Zwaan & Radvansky, 1998): (1) protagonists and objects, (2) causality, (3) intentionally, (4) temporality, and (5) spatiality. Hence, the representation here consists of a protagonist 'John', engaged in an action (intentionality) of 'drinking', the object of this action is 'coffee', and the location is the 'sitting room'.

An important consideration for the Integration stage is the information that is kept in focus in the representation throughout comprehension. The *Landscape Model* (van den Broek, Young, Tzeng & Linderholm, 1999) hypothesises that during reading, constraints on working memory mean that concepts are constantly fluctuating in activation levels. Autonomous activation, as well as higher-level processes involved in searching for links among concepts and events, in turn creates a diverse 'landscape' of activation. This phenomenon will obviously affect the ease with which incoming events are integrated into the representation. While at present our model does not take into account these complex activation processes, it does recognise three distinct types of representation: the *current representation* (events in the currently processed sentence), the *past representation* (events that have previously occurred) and the *implicit representation* (knowledge that is strongly associated with the concepts mentioned in the scenario, as governed by the KB). For example, after reading the first sentence in Table 1, the implicit representation might contain the fact that *"coffee is often contained in a cup"*, and that *"a sitting room is a room in a house where people can relax"*. This information, while not in focus at the time of comprehension, may be called upon later if required when attempting to integrate future events.

Once the initial representation has been constructed from the events in the first sentence, the model can now process the next. As before, this sentence must be checked for coherence in the KB, but in addition to making sense in isolation, the events in this sentence must also be shown to make sense *in context*. To do this, the model attempts to link all the concepts in the current sentence with those that

---

[1] It should be noted that in WordNet, there are a number of senses for each word. For example, *'coffee'* can be defined as a beverage, a plant, a seed or a colour. We only included the relevant senses for the component words in the KB (e.g. in Table 1, coffee is a beverage). However, possible expansions of the model might automatically select the correct sense based on the story context.

have previously been mentioned, using information in the current, past and implicit representations. For example, in the second sentence of the predictable condition in Table 1, *"The cup balanced on the armchair",* the model can link *cup* with *coffee* in the previous sentence, as the definition (or gloss) of cup includes the fact that it can be used to contain a beverage. This also makes use of the hierarchical structure of the concepts in WordNet (i.e., the model must know that 'coffee' is a beverage in order to relate it to 'cup'). Such a linking mechanism can be said to constitute a backward inference, or a causal-bridging inference, and is extremely prevalent in reading (Graesser et al, 1994).

Following this, the model creates an *incoherency score* for the sentence in question. This is simply calculated by dividing the number of concepts that have not been successfully linked with the total number of concepts in the sentence. Based on this, if the second sentence is deemed sufficiently coherent, the model updates the representation by integrating the new event(s) in. These events will now be currently in focus, but prior events will still be present in the representation (e.g., in this case, the knowledge that John is the protagonist and that the location is the sitting room). Subsequent sentences in the scenario are processed in the same way, in so far as they are verified in the KB, checked for coherent links with the past representation and incorporated into the current representation.

At the end of the Integration stage, the model will have a *total incoherency score* for the scenario, which corresponds to the overall ratio of unlinked concepts to linked concepts. Conceptually this refers to the amount of 'new' information that could not be inferred by the story context. This will in turn be relevant when making the surprise rating for the critical event, which is governed by the Analysis stage.

Analysis stage

In the Analysis stage, the model outputs a judgement of surprise for the final sentence in the scenario. In a nutshell, it ascertains how well the events in this sentence can be supported by the prior discourse. To accomplish this, the model employs similar linking mechanisms to those in the Integration stage; however it is conceived of as a more effortful search to detect any supportive information for the constituent concepts. There are two main types of link involved here. The first are *direct links* (*D*): these occur when a concept in the final sentence has previously been mentioned in the discourse. If this is the case, the concept is present in the prior representation, and thus should be very easy to re-integrate (van den Broek et al., 1999). For example, in the final sentence of Table 1, *"The coffee spilt all over the carpet"*, *coffee* has already been mentioned in the first sentence so should be very unsurprising. A direct link like this is assigned the highest score of 1. The second type of link is termed an *indirect link* (*I*). Although a certain concept may not have been mentioned before, it could have been implied by the prior discourse, or could be easily inferred based on background knowledge. Here the contents of the implicit representation may come into play. For

example, in the above sentence, *carpet* can be indirectly linked to *sitting room* in the first sentence as sitting rooms often contain carpets. An indirect link is arbitrarily set as half that of a direct link in order to reflect its diminished importance. However we have also varied the relative contributions of the weights of this value, as will be seen in the Simulation section.

As Figure 2 illustrates, the surprise rating is calculated by dividing the sum of the linked concepts (*D* + [*I**0.5]) with the number of concepts in the final sentence (*W*), and getting the inverse of this score. This is then added to the total incoherency score (*IC*) as obtained in the Integration stage, so as to take into account how well the entire representation fits together when making the surprise rating

$$Surprise = IC + \left[1 - \frac{D + (.5)(I)}{W}\right]$$

Figure 2: Surprise rating formula.

This formula allows the model to detect subtle differences in surprise level across the experimental scenarios. For example, the predictable version from Table 1 would be judged as very unsurprising because all the concepts in the final sentence can be easily inferred from the preceding discourse. In the neutral and unpredictable versions, the key concept of 'spill' cannot be causally inferred making these versions appear more surprising than the predictable one. However, because there is a greater overall coherency in the neutral scenario (as governed by the total incoherency score), this will be perceived as slightly less surprising than the unpredictable scenario.

Model Simulation and Evaluation

The performance of the model was tested against that of the participants from the first experiment carried out by Grimes-Maguire and Keane (2005a). For this purpose, we carried out simulations on a number of the scenarios read by participants in this experiment. We then analysed the individual contributions of the different variables required to make the surprise rating by means of a sensitivity analysis. Though based on a relatively small set of data, the model is designed to be generalisable to different types of texts so as to assess surprise for a wide variety of events.

Method

Materials Nine of the 18 scenarios used in Grimes-Maguire and Keane (2005a; exp. 1) were used in the simulation. Each of these was five sentences long and had three different conditions of predictable, neutral and unpredictable (as in Table 1). All 27 scenarios were translated into propositional format for the purposes of the experiment.

Procedure The model took as input each scenario sentence-by-sentence and then outputted a surprise rating using the formula in Figure 2. These scores were then standardised and translated into a number between 1 and 7 (with 1

referring to low surprise and 7 referring to a high level of surprise), to allow for direct comparison with the participant generated scores of Grimes-Maguire and Keane (2005a).

Results & Discussion

In sum, the results of the simulations corresponded strongly with the surprise ratings given by participants. There was a good correlation between the model's scores and the experimental data for the same materials (*Pearson's r* = 0.8, *p*<0.001, *N* = 27). A scatterplot illustrating this correlation can be seen in Figure 3. A regression analysis subsequently confirmed that the model could be used to predict people's surprise ratings for the scenarios ($r^2 = 0.64$, p <0.001).
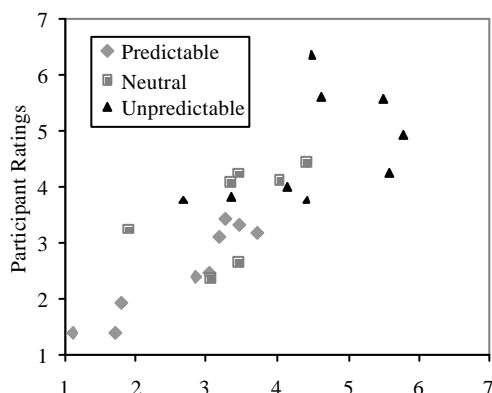
Figure 3: Scatterplot illustrating correlation between model and participant generated surprise ratings

We wished to see how the model would perform in relation to the three conditions of predictable, neutral and unpredictable. Accordingly, we performed a one-way ANOVA, repeated measures, which revealed a significant effect of condition, $F(2,24) = 7.073$, $p < 0.001$, $MSe = 1.075$. As expected, the model rated the predictable scenarios as the least surprising ($M = 2.685$, $SD = .915$), followed by the neutral ($M = 3.826$, $SD = 1.143$) and unpredictable scenarios ($M = 4.503$, $SD = 1.039$). All these conditions differed significantly using Bonferroni adjustments (all *ps* < 0.001). This compares favourably to the experimental results (see earlier Table 1) and suggests that the model was able to detect the varying levels of support or enabling events for the final sentence, as afforded by the different scenario versions.

Assessing contribution of key parameters As Figure 2 illustrates, we had assigned certain values to the parameters required for making the surprise rating *a priori*: namely, any direct links were assigned a value of 1 (100%) and indirect links a value of 0.5 (50%). We also attached a weighting to the total incoherency score (100%). However, we wished to investigate the varying contributions of each of these parameters, and so consequently performed a sensitivity analysis on the data. Many researchers (e.g. Connell & Keane, 2006) have shown that this is an effective technique for assessing the robustness of a model.

Firstly, we carried out a multiple regression analysis to determine the relative contribution of each variable to the power of the model. Thus, total incoherency (*IC*), direct (*D*), and indirect (*I*) links were used as predictor variables, with the surprise rating as the criterion variable. The standardised regression weights from this analysis were .826 (total incoherency), -.305 (direct links), and -.285 (indirect links), all *ps* <0.0001. This illustrates that all three variables contribute to the predictive accuracy of the model. We also performed correlations between each of these variables. As can be seen from Table 3, surprise is highly correlated with total incoherency, which suggests that the Integration stage is very important in determining the surprise rating. Indirect links are also strongly correlated, however direct links do not have a significant relationship with surprise level. This is probably due to the fact that there was little difference in the number of direct links across the three conditions in the scenarios employed.

Table 3: Correlations between variables used in the model

	Total IC	**Direct**	**Indirect**
Surprise	.930**	-.105	-.632**
Total IC		.147	-.523**
Direct			-.276

Next we systematically varied the weights of these contributing variables to ascertain the robustness of the model. As only a weak correlation between direct links and surprise rating was observed, we chose to focus this analysis on the other two measures. Table 4 displays the resulting correlations when varying the total incoherency score (0-100%) and the indirect links (0-100%). As can be seen, when neither indirect links nor the incoherency of the scenario are taken into account, the correlations are not reliable, while increasing the weight attached to both these variables augments the significance. Using the original two variables of indirect (50%) and direct links (100%), we can see that the model performs best when total incoherency score is weighted at 50%. This might suggest that in our

Table 4: Sensitivity analysis for variables total incoherency score and indirect links (direct links held at 100%)

Total IC	*Indirect Links*										
	0%	**10%**	**20%**	**30%**	**40%**	**50%**	**60%**	**70%**	**80%**	**90%**	**100%**
0%	0.2637	0.3974	0.5229	0.6235	0.6918	0.7311	0.7497	0.7554	0.7538	0.7485	0.7413
25%	0.5603	0.6586	0.7322	0.7798	0.8058	0.8164	0.8171	0.8121	0.8040	0.7944	0.7843
50%	0.6922	0.7471	0.7853	0.8089	0.8213	0.8257	0.8246	0.8201	0.8135	0.8057	0.7974
75%	0.7317	0.7645	0.7875	0.8025	0.8110	0.8147	0.8148	0.8125	0.8084	0.8033	0.7974
100%	0.7398	0.7619	0.7781	0.7893	0.7964	0.8003	0.8016	0.8010	0.7990	0.7960	0.7922

original formula we may have been attaching too much significance to this factor, however such a minor shift in correlations does not appreciably affect the power of the model. Also, it would be inappropriate to over-fit this model based on such a small sample.

General Discussion

In this paper we have presented a novel theory and model of surprise judgements which holds that a person's level of surprise at a given event is based on how well that event can be integrated into their discourse representation. The computational implementation of this model has yielded promising results, in that the surprise ratings generated for a number of short scenarios closely mirror participant responses. We have also demonstrated that the principle variables involved in both the Integration and the Analysis stages are important in the assessment of surprise.

Many existing theories of surprise define it in terms of expectancy-disconfirmation, or schema discrepancy (e.g. Meyer et al, 1997; Teigen & Keren, 2003). However, the present work has revealed that the theory of Representation-Fit can offer a more comprehensive account of surprise. The implementation of this theory explains the ability that people have to distinguish different levels of predictability in discourse scenarios; it suggests that they can search for and detect the strength of enabling factors for any given event in a depicted situation. Consequently, surprise does not only occur following unexpected events, rather it is a more complex assessment involving both automatic and strategic processes. It is important to note however that we do not suggest expectation is not involved in surprise at all, rather we propose this is not the *only* factor in the phenomenon. Recent empirical work by Maguire and Keane (2006) lends additional support to this claim, illustrating that, even when events go against expectations, surprise can be lowered if participants have a means of integrating the unexpected event into their representation.

Clearly, there are a number of possible extensions for this model. One option, for instance, would be to further acknowledge the fluctuating levels of activation among the constituent concepts (as in the Landscape Model, van den Broek et al, 1999), and place more emphasis on constraints such as working memory and attentional capacity. However, while this model involves processes important to reading, it is not intended to offer a detailed account of discourse comprehension. Instead it is designed to be generalisable to a number of areas relating to surprise. The model might be used, for example, to explain incidences of surprise in everyday life, or how people use surprise to reason about the likelihood of future events.

In conclusion, this work has shown that surprise is strongly governed by how well events can be integrated into the reader's representation. These initial simulations are promising, and open a lot of areas for future research. As well as shedding light on the phenomenon of surprise, the results illustrate the complex nature of event representation and discourse comprehension.

References

Black, A., Freeman, P. & Johnson-Laird, P.N. (1986). Plausibility and the comprehension of text. *British Journal of Psychology, 77,* 51-62.

Connell, L. & Keane, M.T. (2006). A model of plausibility. *Cognitive Science, 30(1),* 95-120.

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

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

Graesser, A. C., Singer, M. & Trabasso, T. (1994). Constructing inferences during narrative text comprehension. *Psychological Review, 101*, 371-395.

Grimes-Maguire, R. & Keane, M.T. (2005a). Expecting a surprise? The effect of expectations in perceived surprise in stories. *Proceedings of the 27th Annual Conference of the Cognitive Science Society.* Hillsdale, NJ: Erlbaum.

Grimes-Maguire, R. & Keane, M.T. (2005b). A cognitive model of surprise in narrative. *Kogwis: Proceedings of the Seventh Biannual Meeting of the German Cognitive Science Society.*

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.

Maguire, R. & Keane, M.T. (2006). Surprise: Disconfirmed expectations or representation-fit? *Proceedings of the 28th Annual Conference of the Cognitive Science Society,* Hillsdale, NJ: Erlbaum.

Miller, G. A. (1995). WordNet: A lexical database for English. *Communications of the ACM, 38(11).*

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.

Shackle, G.L.S (1969). *Decisions, order and time in human affairs.* Cambridge: Cambridge University Press.

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

van den Broek, P., Young, M., Tzeng, Y., & Linderholm, T. (1999). The landscape model of reading: Inferences and the on-line construction of a memory representation. In H. van Oostendorp & S. R. Goldman (Eds.), *The construction of mental representations during reading* (pp. 71-98). Mahwah, NJ: Erlbaum

van Dijk, T. A., & W. Kintsch. (1983). *Strategies of Discourse Comprehension.* New York: Academic.

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