

Saillenfest, Antoine & Dessalles, Jean-Louis (2014). **Can Believable Characters Act Unexpectedly?**, *Literary and Linguistic Computing* (2014) 29(4): 606-620. doi: 10.1093/lc/fqu042

Can Believable Characters Act Unexpectedly?

Antoine Saillenfest

Telecom ParisTech, France

Jean-Louis Dessalles

Telecom ParisTech, France

Correspondence to:

Antoine Saillenfest

INFRES, Department of Computer Science and Networks,

Telecom ParisTech,

46, Rue Barrault,

75634 Paris Cedex 13,

France

Email: antoine.saillenfest@telecom-paristech.fr

Phone: + 33 145813119

Can Believable Characters Act Unexpectedly?

Abstract:

Unexpectedness is a major factor controlling interest in narratives. Emotions, for instance, are felt intensely if they are associated with unexpected events. The problem with generating unexpected situations is that either characters, or the whole story, are at risk of being no longer believable. This issue is one of the main problems that make story design a hard task. Writers face it on a case by case basis. The automatic generation of interesting stories requires formal criteria to decide to what extent a given situation is unexpected and to what extent actions are kept believable. This paper proposes such formal criteria and makes suggestions concerning their use in story generation systems.

1. The unexpectedness-believability dilemma

Interest in narratives crucially relies on the author's ability to design unexpected situations. The exercise requires however a bit of caution, as the following example illustrates.

***The Knife Story.** John and Mary are true lovers. Over time, their love is growing. On that Tuesday, Mary has breakfast with John as usual. She stands up, goes to the kitchen, grabs a knife, returns and stabs John in the back.*

This example illustrates the mutual exclusion between unexpectedness and believability. Mary's action does not make sense at this point of the story. Readers are in a desperate need for an explanation that will restore Mary's rationality. In the absence of such explanation, or if it comes too late, Mary's character is at risk of appearing non believable, and the story's attractiveness will suffer as a result.

Few authors addressed the issue on a general basis. The narrative generator *Prevoyant* (Bae and Young, 2008) is an attempt to generate flashback and foreshadowing, specifically targeted at the evocation of surprise in the reader's mind. Surprise is mainly aroused by the manipulation of temporal structures in the narrative. Foreshadowing provides the reader with expectations and flashback provides the reader with an explanation of the surprising event. Using a reader model, the system evaluates both the presence of an unexpected event and the fact that the story structure as a whole will hang together and make sense to the reader. *Suspenser* (Cheong and Young, 2008) is a framework that determines narrative contents intended to arouse high level of suspense in the reader. This system relies on the idea that a reader's suspense level is affected by the problems that characters must face and by the number of solutions available to them. The system manipulates the story events in order to increase or decrease the chances of success and also proposes a measure of the level of suspense. Both studies, *Prevoyant* and *Suspenser*, address the question of generating coherent stories that arouse an effect in the reader. However, the problem of creating interesting

situations by generating surprise or suspense is addressed considering only some specific aspect of these notions. No general theoretical framework that would serve as guidance for the generation of interesting and believable narratives is provided. Reconciling unexpectedness and believability is generally considered as a skill that often makes the difference between good and average quality fiction. The automatic generation of interesting stories requires some formal criterion to decide to what extent a given situation is unexpected and whether actions are kept believable. The aim of this paper is to offer such a formal criterion and to show how it could be implemented in story generation systems.

In what follows, we will first illustrate with an example (the Grail Story) how unexpectedness is used by authors to raise interest. They often bring characters to the frontiers of believability. To reproduce this skill, a theoretical framework is needed. We consider such a framework, *Simplicity Theory*, in a further section. After a brief presentation of the theory, we will show how it can be used to formally characterize two notions: surprise and intentionality. Thanks to these definitions, we will be able not only to express the unexpectedness-believability dilemma, but also to get around it. We will then go further, showing how the theoretical definitions can lead to practical calculus and, eventually, could be used in story generation systems.

2. The Grail Story

Humans are narrative beings. They spontaneously tell dozens of narratives on a daily basis about events that occurred recently or that are related to previously mentioned events (Tannen, 1984, Norrick, 2000). Telling stories can be useful for several functions, including learning and remembering (Renninger and Hidi, 2002; Hidi and Renninger, 2006), advertising and, of course, entertaining (Kintsch, 1980). For all these purposes, raising interest seems to be a major requirement for engagement. Conversational stories must be interesting to the audience, or else narrators are exposed to being perceived as socially inept (Polanyi, 1979). A few studies have addressed the question of what makes a story interesting and several factors have been listed. One

of these factors is the reference to major life themes that elicit strong affective reactions, such as death, religion, sex or politics (Kintsch, 1980; Schank, 1979). This effect is often called *emotional interest* (Kintsch, 1980). It crucially depends on the audience's personal values, personal experiences, personal emotions or pre-existing knowledge (Schraw and Lehman, 2001). Besides this emotional factor, surprise also seems to play a crucial role in eliciting interest in the narrative context. This notion of surprise is often expressed in terms of *unexpectedness* (Batsch and Estes, 1997). Several studies highlighted the combined role played by affective reactions and unexpectedness in the elicitation of interest (Galtung and Ruge, 1965; Labov, 1972; van Dijk, 1993).

Unexpectedness, however, is not a cheap narrative feature. It carries the risk of affecting the entire coherence of the story and especially the characters' believability, as illustrated with the Knife Story above. Authors are nevertheless able to generate unexpectedness in a way that the audience may accept as believable in the context of the narrative. Let's consider the following story, which is an edited excerpt from a film script, *Indiana Jones and the Last Crusade* (Story by G. Lucas and M. Meyjes, 1989).

The Grail Story.

At that moment DONOVAN and ELSA appear. Donovan turns to face the display of cups.

DONOVAN: *Which one is it?*

KNIGHT: *You must choose. But choose wisely. For as the True Grail will bring you life, the False Grail will take it from you.*

DONOVAN: *I'm not a historian. I have no idea what it looks like. Which one is it?*

ELSA: *Let me choose.*

DONOVAN: *Thank you, Doctor.*

[...]

ELSA chooses a cup—a solid gold, emerald encrusted goblet. DONOVAN instantly takes it from her.

DONOVAN: *Oh, yes. It's more beautiful than I'd ever imagined. This certainly is the cup of the King of Kings.*

DONOVAN rushes to the well and fills the goblet with water.

DONOVAN: *Eternal life!*

DONOVAN drinks from the goblet. Then, DONOVAN'S entire body starts to convulse. His face contorts in agony. He grabs his stomach and turns toward ELSA. [...] He starts to age—fast! [...] As he falls his body breaks into flames, then shatters against the wall.

KNIGHT: *He chose...poorly.*

INDY studies the array of chalices.

ELSA: *It would not be made out of gold.*

INDY picks up another cup—a simple earthenware jug.

INDY: *That's the cup of a carpenter.*

He and ELSA exchange a look.

INDY: *There's only one way to find out.*

INDY goes to the well and fills the earthenware jug with water, then pauses. INDY brings the jug to his lips and takes several large swallows. A strange sensation overcomes him, a feeling of peace and contentment...and we see his wounds begin to heal.

Emotion is of course central in this fragment which deals with extreme outcomes: eternal life or immediate death. However, merely mentioning death is not enough to share emotion. To be interesting, a situation must be out of the ordinary, exceptional, surprising (Labov, 1972:370). Surprise is known to play a crucial role in eliciting interest (Reisenzein, 2000). It seems inherently attached to the unexpectedness of situations. Both adults and children seem to conceptualize surprise in terms of unexpected events (Bartsch and Estes, 1997). We will use the notion of

unexpectedness to capture the common feature of situations that are perceived as surprising, extraordinary or amazing.

One can observe that several elements in the Grail Story concur to make the situation unexpected to some observer. Donovan's death is highly unexpected, at least to him, as far as the choice of the most beautiful chalice is believed to be obviously correct. Indy's choice of the most insignificant chalice is unexpected as well, because the object blatantly contradicts standard expectations about the Holy Grail. Lastly, Indy's survival is unexpected to an observer who is not fully confident that he made the correct choice. Indiana Jones is an adventure film. People familiar with this film genre usually expect the hero to survive till the end and the villain to lose or to die. However, this belief, or even this knowledge (if one has already seen the movie), may be cancelled due to narrative transportation. Narrative transportation occurs when a story receiver loses herself in a story and gets detached from the "real world" because of her engrossment in the story (Gerrig, 1993; Green, Brock and Kaufman, 2004) . Receivers empathize with the characters and, being immersed in the story, share their local expectations. This detachment from the "reality" may explain why, despite having already seen the film or having global expectations associated with the genre of the narrative, one can nevertheless be surprised.

Several elements in the Grail Story illustrate a strange fact about unexpectedness: interest and emotions do not seem to depend on "who" is supposed to experience it. Unexpectedness is generally experienced directly by the receiver. This is what happens when Indiana Jones takes the wooden cup in the Grail story or when Mary stabs John in the Knife story. However, fully expected situations may be thrilling just because they are unexpected to a character. This phenomenon results from identification with the character. Identification goes beyond narrative transportation. Cohen defines it as "a process that consists of increasing loss of self-awareness and its temporary replacement with heightened emotional and cognitive connections with a character" (p. 251). In other words, the story receiver may see the storyworld through a character's perspective. Events that are unexpected to the character become automatically interesting.

3. Unexpectedness, believability and coherence

If, as we claim, unexpectedness controls narrative interest and emotions, it is crucial to define it operationally. A typical view consists in defining unexpectedness by the low probability of events, as there seems to be an inverse linear relationship between surprise and probability (Reisenzein, 2000). However, the picture must be more complex, since not all low probability events are necessarily surprising. Teigen and Keren (2003) suggested that surprise can be determined by the contrast between an outcome and an expectation, rather than by the absolute probability of the outcome *per se*. For example, if you open a jar to take a cookie and you suddenly realize that the box is empty, then you will be surprised because your expectation has been disconfirmed. However, the hypothesis that surprise is a matter of disconfirmed expectations does not take into account situations in which no expectations are formed. For example, if someone enters your room by mistake, the event is surprising to you despite the fact that there was no expectation; the event is unexpected because it has not been expected at all. Ortony and Partridge (1987) describe three causes for surprise: active expectation (or prediction) failure, passive expectation (or assumption) failure, and unanticipated incongruities. Unexpectedness, in a wide sense, must therefore refer not only to low probability events, but also to disconfirmed expectations and to unanticipated events.

A story writer who wants to arouse interest by generating unexpected events may, following Reisenzein (2000), try to create low probability situations. This strategy supposes that probabilities be available. In the Grail Story, there is no way to assess the probability of dying for having chosen the most beautiful cup, nor to measure the probability that a wooden cup may be the Holy Grail. Maguire, Maguire and Keane (2011) propose that the level of unexpectedness experienced should be associated with the difficulty of integrating the unexpected event with an existing representation.

‘For example, if you found your car key was missing, and you had no way of explaining it, then you might experience a high level of surprise. However, if a plausible explanation subsequently emerged that allowed the anomaly to be resolved, such as realising that you left the key inside the car, then the experience of surprise should subside.’ (Maguire, Maguire and Keane, 2011).

The level of unexpectedness therefore depends on the difficulty of accounting for the event in the initial context. This level of difficulty is assessed using our knowledge about the world, including probabilities and causal laws.

Creating unexpected situations may be at the expense of believability. Believability, in a general sense, refers to two aspects: plot believability and characters' believability (Riedl and Young, 2010). Plot believability is associated with the coherence of the storyline. The fictitious world in a narrative has its own physical and causal rules that the author must respect and make the audience aware of. This knowledge about the rules of the world may be partially provided by the genre of the narrative. Indiana Jones adventures belong to the fantastic-adventure genre in which “eternal life” is a possible outcome. However, Indiana Jones lives in a fictitious world that resembles our everyday world. Events such as an earthquake or a death by cardiac arrest are 'rare events', even in that fictitious world. Introducing such unexpected events out of the blue to alter the sequence of events at the appropriate moment, like a *deus ex machina*, may jeopardize the believability of the scenario. If such an event happens, it seems that from now on everything is possible and nothing is unexpected.

Characters' believability relies on rules and coherence as well. One of the main components of the characters' believability is their intentionality. Agents must be intentional, which means that their actions must be intentional and recognized as such. Irrational or incoherent action would negatively impact the audience's suspension of disbelief (Bates, 1994). In this paper, we focus on the recognition of the characters' intentions, as it is the main prerequisite for characters' believability.

In the Knife Story, the believability of the character named Mary is put at risk. Mary's action is highly unexpected, but it is inconsistent with our knowledge of the character's values. To restore Mary's believability, one has to imagine that Mary's action has been performed for some reason that remains still unknown. In acceptable stories, believability, in a general sense, may be put at risk, but only for a while.

In the Grail Story, the characters' actions and their consequences are definitely surprising, but they are immediately kept consistent with our knowledge about the situation. Donovan drinks from the golden cup because he believes it to be the Holy Grail and because he wants eternal life. Indiana Jones drinks from the earthenware cup because he has reasons to believe, up to a point, that it is the true Grail. His choice of the earthenware cup is initially at odds with our prototypic knowledge about the Holy Grail, but the inconsistency lasts only for a few seconds, until Jones explains his choice. This illustrates the importance of restoring logical consistency by providing relevant information that enables the audience to reintegrate the odd event into a coherent context. In the context of narratives, this requirement has been named by Kintsch (1980) *postdictability*.

'The text as a whole must hang together and make sense to the reader, so that he is able to construct a coherent macrostructure in which each text unit has its place and is meaningfully related to other sections of the text. Weird and unpredictable statements in a text are interesting only in so far as they are well motivated within the text as a whole, at least by hindsight.' (Kintsch, 1980:89)

The postdictability constraint is a precondition for believability. In particular, the characters' *intentions* should be manifest to the audience. This is the main source of the dilemma: how can transparent intentions lead to unexpected actions? One possibility is to delay the disclosure of characters' intentions. Actions appear unexpected in the first place, and the audience's interest is aroused by their unexpectedness. Then, in a second phase, some piece of information is revealed that restores the characters' believability. This scheme is used in many narratives, including situations in which the character's destiny is perfectly known in advance. The whole point of a legend like Oedipus myth is to reveal how reasonable circumstances can lead a believable character to perform acts that are at odds with his values (kill his father, marry his mother). The whole plot consists in restoring postdictability by progressively revealing those circumstances.

The problem with postdictability is that there is no procedure to guide its design. In the Knife Story, one may be told afterwards that Mary is a spy that pretended to love John just to murder him. Her act is now less unexpected, but the unexpectedness is now transferred to the spy story. If we

want to generate scenarios automatically, we need a measure of believability and of unexpectedness that help us predict whether the scenario is acceptable. We will use the Simplicity Theory framework to do this.

4. Simplicity Theory

Simplicity Theory (ST) provides formal definitions of unexpectedness and intention. It is therefore particularly relevant to narrative design. ST was initially developed to predict interest in spontaneous conversational (non fictitious) narratives (Dessalles, 2008). Its central claim is that interesting topics are systematically associated with *complexity drops*. This makes the cognitive notion of *complexity* an essential tool to study narrative interest.

Previous studies showed that humans are sensitive to complexity in the Kolmogorov sense, *i.e.* the size (in bits) of the most concise description (Chater, 1999; Chater and Vitányi, 2003). Some authors suggested that interest is associated with variations of complexity through time (Schmidhuber, 2009). Intuitively, it means that a situation is interesting if the observer is momentarily unable to account for it in a simple way. ST makes a more constraining statement. Within the ST framework (Dessalles, 2008; Saillenfest and Dessalles, 2012; see www.simplicitytheory.org), unexpectedness is defined as a contrast between two aspects of the event: its generation complexity and its description complexity. The *complexity-drop rule* states that unexpected (and therefore interesting) events must be less complex to describe than to generate.

$$\begin{array}{r} \text{generation complexity} \\ - \text{description complexity} \\ \hline = \text{unexpectedness} \end{array}$$

Formally, this can be rewritten as:

$$U(s) = C_w(s) - C_d(s).$$

where $U(s)$ is the unexpectedness of a situation s , $C_w(s)$ and $C_d(s)$ are respectively its generation complexity and its description complexity. The word 'complexity' is used here in its technical sense, and means the 'size (in bits) of the most concise description'. Algorithmic Information Theory defines the complexity of an object as the length in bits of the shortest program that can describe this object using any (universal) computer programming language (Li and Vitányi, 2008). The corresponding cognitive notion of complexity refers to the size of the shortest description that a given human observer can achieve at a given moment. With this definition, a periodic pattern such as a-a-a-a-a-a-a is simpler than a non periodic one such as a-u-k-u-h-a-b. The observer can communicate the former using a short procedure (repeat 'a' seven times). On the other hand, the latter sequence requires each letter to be described in turn, making the overall description significantly longer. Of course, if the second sequence happens to be the observer's name, there is a much simpler procedure to communicate it and its complexity is smaller.

In ST, the notion of description complexity (C_d) corresponds to the preceding definition. In the Grail Story, the main scene involves central characters (Donovan, Elsa, Jones). These characters are 'simple', because they come first in the list. If the scene had involved an obscure character instead of Jones, the description complexity of the situation would have been larger (three extra bits would be required if this character comes in eighth position in the list of importance). The choice of the most and the least beautiful chalices in the Knight's collection is not fortuitous. It makes their description minimal. The simplicity of outcomes, death vs. eternal life, also contributes to the simplicity of the situation. To calculate the description complexity, receivers may use any available piece of information that has been provided to them: objects, protagonists, places where situations take place, temporal information (time of the action or reference to past events). In the case of a coincidence between two situations, one can spare bits to describe one of the coincidental situations

by referring to the other one (Dessalles, 2008). In other words, the cost (in bits) of describing the two situations together is smaller than the sum of the two separate descriptions. The more alike the two situations, the larger the complexity drop. Calculating complexities therefore requires to have a model of the storyworld (Zwaan and Radvansky, 1998).

The other term of unexpectedness is *generation complexity* (C_w). It refers to the minimal distortion of the world required to produce the event. It can also be described as the size of the *minimal explanation* of the situation, *i.e.* the minimal set of circumstances that the observer must imagine for the event to happen. In the Grail Story, picking the correct chalice among thirty is more complex than choosing it among only two. Technically, the first choice amounts to five bits, whereas the second one requires the determination of only one bit. When characters are placed in seemingly hopeless situations, the generation complexity of any salutary event is measured by the complexity of the circumstances that may bring to the character's rescue. By characterizing the situation as 'hopeless', one only means that this complexity has a large value. Generation complexity corresponds, when applicable, to the anticipated probability (measured in bits) that the outcome would happen in the future. Generation complexity is, however, a more general notion than probability (Dessalles, 2008).

Unexpectedness, defined as the difference between C_w and C_d , characterizes all events that contradict expectations. This includes as special cases events that are intuitively associated with an impression of low probability. Rare events, remarkable happenings, oddities, exceptions, deviations from norms are all extra-ordinary. They contrast with ordinary situations which, almost by definition, conform to expectations and are regarded as probable. These extraordinary situations, as far as they appear to be exceptions, can be distinguished using a simple description. A giant can be distinguished by merely mentioning that her size is extreme. An average person can be distinguished as well, but only by using a lengthy description.

Unexpectedness, however, cannot be deduced only from probabilities. The probability of Mary's aggression in the Knife Story or the probability of Jones's unexpected choice in the Grail Story is hard or impossible to compute. Moreover, events of equal probability may not arouse equal interest. For instance, a lottery draw like 1–2–3–4–5–6 is much more thrilling than a 'normal' draw like 13–23–24–31–35–44, despite their identical probability. The complexity-drop rule makes correct predictions in a wide range of situations such as coincidences, rare or atypical events and complex causal stories (Dessalles, 2008, Saillenfest and Dessalles, 2012, www.simplicitytheory.org). The lottery draw 1–2–3–4–5–6 is highly unexpected, not because it is objectively improbable, but because it is simpler than expected. A 'normal' draw like 13–23–24–31–35–44 is not unexpected: for most observers, its description complexity is as large as its generation complexity. This explains why, despite being objectively equally probable, the consecutive draw is regarded as incredibly unexpected. Maguire *et al.* (2013) refer to 'randomness deficiency' to characterize this effect. ST captures the phenomenon with the notion of unexpectedness, as defined above.

Unexpectedness is necessarily subjective. It depends on the internal knowledge of who is measuring it. As Teigen & Keren (2003) put it:

an astrophysicist may be extremely "surprised" to learn that the moon is one billion years older than previously thought, whereas a layman may declare himself not surprised at all; not because the new estimate was expected, but because it did not conflict with any previously held beliefs (Teigen & Keren, 2003).

In terms of complexity, the main difference between the astrophysicist and the layperson is that the situation is harder to integrate in the world as the astrophysicist knows it than in the world as the layperson knows it. Generating the event, using the astrophysicist knowledge, is (considerably) more complex than generating it using the layperson's knowledge about the world. In the absence of any knowledge about the moon's age, the new fact can be immediately integrated into the world and its generation complexity amount to zero. The astrophysicist, on the other hand, must revise much previous beliefs about the solar system and imagine a set of causes that make the older age logically possible. In the Knife Story, killing John is unexpected to John and to the audience, because it

conflicts with previous beliefs about Mary's psychology. The minimal set of causes (she got mad, or she has been threatened, or she is jealous, or she is a spy...) comes with a generation cost.

As observed by Schmidhuber (2009), complexity may vary through time for a given observer. The same is true for unexpectedness. An event which appears as unexpected at a given moment may become ordinary a moment later. This phenomenon is essential for postdictability. What seems hard to explain (or generate) at a given time may appear simpler to explain some time later with additional information about the world. For example, if one learns that Mary is a spy and that John is her target, her act is less complex to explain. Even her apparent love for him comes for free with the explanation. Description complexity may also vary through time: a lottery draw may initially seem complex, and thus boring. If your best friend recognizes her mother's phone number in it, then the draw becomes simpler and, automatically, more interesting.

5. Simplicity effects in narratives

Simplicity Theory offers an adequate theoretical framework to formally describe notions such as surprise and believability. We examine these two aspects in turn.

5.1 Surprise and complexity drop

Simplicity Theory covers the various characterizations of surprise that have been mentioned. The complexity-drop rule correctly predicts that low probability events, disconfirmed expectations and unanticipated events will all be perceived as unexpected. Low probability events correspond to situations that the observer considers to be hard to generate in the world as she knows it (C_w large). For instance, individuals are rarely struck by lightning. If being struck by lightning is compared with a uniform lottery with N possible outcomes, then the corresponding probability is $p = 1/N$ which corresponds to a generation complexity $C_w = \log_2(N)$ (as one needs that number of bits to decide among the N outcomes). ST adds the constraint that the event must be simple to describe (C_d small). The lightning event is more interesting if the victim is your neighbor rather than if she is any

person in the city. Designating a neighbor requires a much smaller amount of information than the quantity $\log_2(P)$ required to designate one of the P city inhabitants. Low probability events, which are difficult to produce (N large), are unexpected only if their description is simple (simpler than $\log_2(N)$).

The surprise generated by disconfirmed expectations and unanticipated events cannot be explained by low probability. What would be the probability that the recently filled cookie jar be empty? What would be the probability of Mary's assault against John in the Knife Story? The audience measures the unexpectedness of the situation without relying on statistics concerning emptied jars or stabbing by loving partners. Here also, complexity drop is the answer, and ST provides methods to evaluate the level of unexpectedness of the situation.

ST defines the strength of beliefs by the unexpectedness of their negation (Dessalles, 2008). Beliefs strengths, as complexities, are measured in bits. When you find the jar empty, you rely on the strong belief that you filled it the day before. When this belief is contradicted, its strength gives the amplitude of the event's generation complexity. Subsequently, you may find a causal scenario that will diminish this initial value. The causal scenario may require several independent circumstances c_1, c_2, \dots, c_k to have occurred (the fact that your nephew entered the home, the fact that he knew where the cookies were hiding, the fact that he eat them all, the fact that he left before your return). If this is the simplest scenario you may think of, then C_w is the sum $C_w(c_1) + C_w(c_2) \dots + C_w(c_k)$. And if this value is smaller than your belief that the jar should have been full, then it will replace the initial measure of the event's generation complexity.

Similarly, the initial unexpectedness of Mary's aggression depends on your belief that she would not perform such an act. This belief is closely related to the recognition of the character's intentions. This issue will be addressed in the next section.

A naive implementation of the complexity-drop rule would consist in generating maximally unexpected outcomes. For instance, the hero of the story may die at the most inappropriate moment

from cardiac arrest. This works perfectly well in real life. For instance, John Paul I's death just 33 days after his election as pope raised considerable emotion, well beyond the catholic world. The technique wouldn't work in a work of fiction. One may imagine a situation in which Donovan dies from cardiac arrest just at the right moment, as he is threatening the hero's life. This would certainly be unexpected, but also not believable. There is an important difference between an event that happens 'in the real world' and an event that happens in the context of a fictitious narrative. The alternative reality of fictitious worlds has its own rules and constraints. In Indiana Jones' world, eternal life is possible, but spontaneous death is still the result from a lottery. Donovan may indeed die from cardiac arrest, but generating his death at the right minute requires a lot of information (some 24 bits if the moment of his death is determined uniformly over 40 years; this corresponds to the complexity of generating 24 times Head in a row while flipping a coin). The audience may accept the alternative world as it is initially presented, but not the extra information used to generate a particular event. It is the author's primary duty to keep the fictitious world believable to the audience. The audience must know enough about the rules of the fictitious world and about its state in order to anticipate what may occur and what may not occur. Usually, these rules are presented by the author to the audience in the novel. The genre of the story may also provide information about what is possible or not, due to some conventions or norms. Though Indiana Jones lives in a fictitious world that resembles our everyday world, it is assumed that in this kind of adventure/fantastic movies, objects may have mythological powers such as destroying one's body instantly or providing eternal life. Without the ability to anticipate what may occur and what may not, everything is regarded as possible. As a consequence, generation complexity C_w drops to insignificant values and the audience is never surprised. Similarly, if the world has constraints but the author violates them too easily, these constraints tend to be considered worthless and the consequence is the same: loss of believability, low C_w , no surprise. There is therefore an inherent contradiction that emerges from the very notion of unexpectedness used in a fictitious context.

Events that are too difficult to integrate in the fictitious world put it at risk of appearing unbelievable.

5.2 *Intentional characters*

Believability is a necessary property, not only of the fictitious world, but also of characters. Believable characters must have coherent intentions. They are expected to make attempts to solve their problems, and conversely all their actions must have purposes. In the Knife Story, the purpose of Mary's act is opaque. The character is at risk of losing believability until this purpose is provided (e.g. she is a spy who pretended to love John). As Riedl and Young (2005) put it: 'one important aspect of character believability is character intentionality [*i.e.*] the way in which the choice of actions and behaviors that a character makes appears natural (and possibly rational) to external observers.' We will consider here that agents are believable if they are *intentional*.

According to classical models, characters behaving intentionally must have a desire (for an outcome) and appropriate beliefs (about how their action would lead to that outcome) (Forguson, 1989). This belief/desire model is however incomplete. Shaver (1985) proposed a more precise definition: an action is intentional if the agent has a desire for an outcome, beliefs about the consequences of that action, and beliefs about his or her ability to perform the action. More recently, Malle and Knobe (1997) have listed five necessary components to determine if a given outcome has been intentionally achieved. These components are the desire for the outcome, the beliefs that the action would lead to the outcome, the intention to perform the action, the awareness of the act while performing it, and a sufficient degree of skill to perform the action reliably.

In the computational field, the BDI framework (Bratman, 1987) represents an agent's behavior as resulting from the interaction between its beliefs, its desires and its intentions. Beliefs are propositions that the agent believes to be true about the world in which it is situated. Desires are states of the world that the agent wishes to make true. Intentions are the desires that an agent

commits to make true. The act of committing to an intention is the process of choosing a set of consistent desires and forming a goal to make those desires true. Once a goal is established, the agent constructs a plan to achieve it. Acting intentionally in order to reach the goal consists in achieving that plan.

These models highlight the fact that to perceive intention, one must recognize that the acting agent had anticipated a causal scenario leading from the action to its outcomes. The reader of a story should therefore be given sufficient information about the character's desires, knowledge and beliefs (about the environment and about his/her own skills). Note that this information needs not be given fully in advance. Authors often provide only partial information to readers about what their characters know. Their actions appear unexpected in the first place and yet intentional once the missing information is disclosed.

Simplicity Theory offers a formal definition of intentionality (Saillenfest and Dessalles, 2013). First, we can observe that the recognition of intentions and the notion of unexpectedness are closely related notions. Among all the actions that can be imagined to reach a desired outcome, an intentional agent should choose the one that minimizes the unexpectedness of the outcome. The intentionality attached to an action is the difference between the emotional intensity attached to the outcome and the unexpectedness of the causal link.

emotional intensity of the outcome

- causal unexpectedness

= intentionality

Formally, this can be rewritten as:

$$Int(s, a) = E(s) - U(s||a).$$

where $Int(s, a)$ (if positive) is the level of intention to make s happen while performing action a , $E(s)$ is the anticipated emotional intensity of the outcome and $U(s||a)$ is the unexpectedness of the situation s once a has been performed. The emotional intensity evaluates the desire (resp. the aversion) one has towards a desired (resp. non-desired) situation. It may be controlled by some parameter, such as the amount of money that may be earned or lost. The intentionality rule states that actions are performed intentionally if the outcome is less unexpected once the action has been performed.

This intentionality rule is also valid for negative intentionality, when one does *not* want to perform an action that would lead to a negative outcome. When an action has several outcomes, the corresponding partial intentionality values add up, with their positive or negative signs, to give the overall intentionality. Typical examples of this situation are moral dilemma, in which an action has both a desired outcome, *e.g.* avoiding a disaster, and an undesired one, such as provoking an innocent's death (Saillenfest and Dessalles, 2013).

The above definition of intentionality captures the fact that intentional agents should perform actions that will expectedly generate desired outcomes or prevent undesired outcomes. The corresponding computation is supposed to be performed by the agent, *i.e.* the character, in a narrative. What does this mean? Characters do not think by themselves, of course. The definition presupposes several levels of cognitive computation. At a first level, the emotional intensity of the outcome is supposed to be anticipated by the acting characters. But this anticipation is evaluated at a second level by the audience. Moreover, the characters' intentionality depends on the audience's readiness to attribute to them sufficient computing power, so that they can anticipate the consequences of their actions. This explains why dogs or children are less likely to be considered as intentional agents than adult humans when performing certain actions. In many stories, higher levels of cognitive computation are involved, *e.g.* when characters are supposed to understand the

intentionality of other characters. For instance, each character in the Grail Story is supposed to understand Donovan's action. This is also true for the knight, despite his knowledge that the cup chosen by Donovan is the wrong one. Lastly, one more level of cognitive computation must be considered if one remembers that the author must anticipate all other levels.

6. Reconciling unexpectedness with believability

When designing a story episode involving an action and its consequences with the aim of making the episode interesting, an author must achieve two objectives. The first one is to maximize the overall unexpectedness of the episode, from the audience's perspective. The second one is to maintain the narrative world's believability and the characters' believability, making characters intentional. Let us observe how the design of the Grail Story serves these two purposes.

Donovan turns to face the display of cups. Donovan must decide which cup, among the 30 cups in front of him, he should choose to get eternal life (or immediate death). The complexity of his choice amounts to 5 bits (as $\log_2(30) \approx 5$). According to ST's definition of the intentionality, the intentionality of Donovan's choice (a_1) results from two contributions. A negative term, with absolute value $E(s_1) - U(s_1|a_1)$, inherited from the perspective of an immediate death (s_1); and a positive term $E(s_2) - U(s_2|a_1)$ inherited from the perspective of reaching eternal life (s_2). The intentionality of a_1 depends on the fact that $E(s_2) - U(s_2|a_1)$ is larger than $E(s_1) - U(s_1|a_1)$. The causal unexpectedness of achieving s_2 is $U(s_2|a_1) = 5$ bits, as the correct choice (the Holy Grail) requires $C_w = 5$ bits to be generated by chance, it is unique ($C_d = 0$) and once it has been made, s_2 becomes certain. $U(s_1|a_1)$ is negligible in comparison. A rational agent should perform a_1 only if $E(s_1) < E(s_2) - 5$, *i.e.* if the agent values eternal life significantly more than anticipated death (note that this lottery test can be used to measure $E(s_2) - E(s_1)$ for any pair of alternatives).

It's more beautiful than I'd ever imagined. The choice of the most beautiful item among the 30 cups makes its description complexity (C_d) close to zero: the object is unique by this obvious

feature. If Elsa, for good reasons, had picked a randomly looking grail, the scene would have been less interesting. The unexpectedness of the actual choice is $C_w - C_d \approx 5$ bits.

This certainly is the cup of the King of Kings. This statement reveals another complexity drop in Donovan's mind. His confidence in the fact that the cup chosen by Elsa is, beyond any doubt, the correct one, and the underlying explanation that the Holy Grail must be the most beautiful one, bring $C_w(a_1)$ from 5 bits to zero. This has an effect on the intensity of his intention: $U(s_2||a_1) \approx 0$ (certainty) and $U(s_1||a_1)$ becomes very large.

Donovan rushes to the well and fills the goblet with water. Suppose that the two perspectives be equally intense: $E(s_1) \approx E(s_2)$. If drinking from the chosen chalice (a_2) has significantly positive intentionality, it must come from the much lower causal unexpectedness: $U(s_2||a_2) \ll U(s_1||a_2)$.

Donovan's entire body starts to convulse. For any reader taking Donovan's perspective, immediate death is mostly unexpected at this point ($U(s_1||a_2)$ large), and its occurrence generates surprise.

Indy picks up another cup – a simple earthenware jug. The choice of the ugliest cup contributes to description simplicity. Designating the last item in the list requires less information than designating an intermediary one. One more time, picking one of the simplest items generates immediate unexpectedness. Moreover, Indy's choice seems to ruin any hope to reach eternal life. The unexpectedness that his choice a_3 be correct was anticipated to amount to $U(s_2||a_3) = 5$ bits, assuming that he would make a random decision. His irrational choice brings $U(s_2||a_3)$ to an extreme value and conversely $U(s_1||a_3)$ close to zero. ST measures $U(a_3)$ by the absolute value of its (very negative) necessity (Dessalles, 2008), which amounts in this case to $U(a_3) = Int(s_1, a_3) \approx E(s_1)$. a_3 is unexpected, because it was the last thing to do. For a moment, Jones ceases to be believable.

That's the cup of a carpenter. This piece of information provides an explanation for John's action. We observe a complexity drop, as $C_w(a_3)$ goes down from an extreme value to a negligible

one. This restores Jones' believability. The picture is inverted: $U(s_2|a_3)$ is believed to be negligible in Jones' mind and his action now appears intentional.

There's only one way to find out. By comparing the decision with a lottery, the author suggests that $U(s_2|a_3)$ is non zero after all.

Indy fills the jug with water, then pauses. The hesitation is not fortuitous. It suggests a lower intention of drinking than in Donovan's case. This lower intention reveals one more time that $U(s_2|a_3)$ keeps a residual value in Jones' mind. This illustrates how an author can use behavior as clue to the character's mental state. Jones uncertainty about the outcome, once shared by the audience, makes the desired outcome more unexpected (and the undesired outcome more expected).

A feeling of peace and contentment. Even if the residual unexpectedness $U(s_2|a_3)$ has a limited value, it represents the actual unexpectedness attached to the realization of s_2 .

The preceding description focuses on unexpectedness and not on emotion. However, unexpectedness is the amplifying factor of emotion (Dessalles, 2008). The extreme stakes (eternal life vs. immediate death) make its role all the more crucial in this story.

The Grail Story illustrates a two-phase mechanism that is widely used in narratives. Let's call it the *up-and-down* mechanism. It consists in introducing a delay between an action and the moment when some piece of information that is crucial for measuring a character's intentionality is delivered. Jones thinks that the Holy Grail cannot be a magnificent chalice, but this element is opaque to the audience when he picks the wooden chalice. The reason of his choice is provided immediately after. However, the disclosure of crucial information may be delayed a much longer time. The audience has to suspend disbelief for a while and must accept that something remains unexplained. The reader of the Knife Story is in desperate need for an explanation for Mary's act. The effect of the explanation (e.g. the fact that she is a spy) is to restore the intentionality of her act. The up-and-down mechanism creates two moments in which complexity drop occurs: a first time

because performing an apparently irrational action is unexpected; and a second time when the relevant information is unveiled and restores the character's rationality. This information simplifies the explanation of the character's behavior (Mary has not gone suddenly crazy after all), and is therefore experienced as a second complexity drop.

The first phase of the up-and-down mechanism leaves the audience powerless when it comes to measuring the unexpectedness of the (seemingly) irrational act. How can absurdity be quantified? The answer comes from the assumption that the character must remain believable. Mary must have some imperious reason for stabbing her lover (a). The intentionality of her act seems very negative, with absolute value $Int(s, a) = E(s) - U(s||a)$, where $E(s)$ is the emotional intensity that she will experience if he dies. This negative intentionality corresponds to the strength of the audience's belief that Mary will not stab her lover. From the audience perspective, $U(a) = E(s) - U(s||a)$. This provides an upper bound to the unexpectedness of Mary's act.

7. Discussion and perspectives

Our main objective in this paper has been to propose operational definitions of two notions that are crucial for the determination of interest: unexpectedness and believability. We could provide formal definitions using the general framework given by Simplicity Theory. We showed that at face value, these definitions lead to an apparent incompatibility between interest and believability. However, authors can take advantage of the fact that the audience's limited knowledge may lead to wrong assessments of intentionality. In particular, what we called the up-and-down mechanism consists in withholding some piece of knowledge that would explain the character's behavior. This behavior is wrongly regarded, during a time, as seemingly irrational and thus as unexpected. In a second phase, the character recovers her believability. This phase of disclosure constitutes a second occasion of complexity drop that may be as pleasurable as the initial one.

Our theoretical framework is based on notions such as cognitive complexity, which is closely related to the mathematical notion of Kolmogorov complexity. A typical objection against the use of that notion is that it is theoretically not computable. This objection does not obtain in the case of cognitive complexity. This notion refers to the most concise description that a given human observer is able to achieve at a given moment. The definition is operational. Its reliance on bounded computational resources and on time is crucial for the up-and-down mechanism to produce its effects. We showed through different examples how the computation of complexity can be achieved in practice (see also www.simplicitytheory.org).

The up-and-down mechanism implements the postdictability principle (Kintsch, 1980). It offers a way of reconciling unexpectedness with believability. A good part of the authoring skill consists in presenting events that will initially seem hard to explain, as hard as suspension of disbelief permits. Then, in a second phase, the author unveils outcomes that are both simple to generate and hard to foresee.

Our model of unexpectedness and intentionality is not just a theoretical tool to account for the cognitive process underlying readers' experience when reading a narrative. It also offers a technical tool that can be used in practice to generate narrative automatically. It is a first step that may help to bridge the gap between abstract notions and a generation process. The principles developed in this paper could possibly be implemented in a computer program that may help scriptwriters in making the determining factors of interest explicit. The tool could also be used for the automatic design of small interesting scenarios in simple worlds.

The most basic principle that can be implemented is the maximization of unexpectedness. The system may choose or modify situations that are (1) emotional, (2) hard to produce (but not too hard to remain believable) and (3) simple to describe. The system may be given control over a variety of parameters: choice of actions, choice of timing, choice of places, choice of characters involved, and so on. To keep the system simple, we suggest using the same planning algorithm at three different

levels. Intentional characters should plan their actions using their own knowledge to achieve desired outcomes. The system should plan the characteristics of events to maximize their unexpectedness. And lastly, the system should plan the global scenario to choose the events composing it.

The three planning phases cannot be independent. By making characters highly believable, the system is at risk of creating a boring scenario, as outcomes will be conceived to be maximally expected. The problem is to design unexpected situations in which each character continues to behave intentionally. To achieve this result, the system may use various tricks, such as introducing conflicting desires or keeping some characters uninformed about certain crucial facts.

One of the main difficulties is to build a causal representation of the world that the system will use to compute the generation complexity of situations. The fictitious world is ideally represented using a (quantitative or qualitative) physical model for things, and a psychological model for characters. A rough approximation consists in representing facts and actions with predicates and linking them with causal and logical rules. This kind of representation is commonly used in agent-based modeling (Fikes and Nilsson, 1971; Riedl and Young, 2010). It is appropriate to deal with abstract aspects of a story. It is less adequate to represent concrete aspects such as spatial relationships.

Among the predicates that represent the fictitious world, Actions have a particular status because they can be ‘performed’, and once performed, the state of the world has been altered. They come with enabling preconditions and a set of resulting facts that become true once the action has been performed. Actions are not the only way to make the world change. We suggest to introduce ‘world events’ that can be defined as world’s actions. World events cannot be controlled by characters. They refer to events that happen in the world and may interfere with anyone’s plans.

The planning abilities of the characters and of the audience can be modulated in two ways. First, they may have different knowledge about the state of the world and about (other) characters’ desires. The output of their planning will be different, in a way that leaves room for unexpectedness.

For instance, in the Grail Story, Jones has one piece of knowledge that both Donovan and the audience lack, namely the fact that the Holy Grail should be a modest container rather than a golden chalice. His plan leads to an action: pick a wooden cup, that differs from Donovan's choice and that appears totally unexpected to the audience. The second way in which planning may differ depending on who is supposed to perform it is due to limits in the inference power. The automated author may count on the audience's cognitive habits or laziness to draw false conclusions. People are used to regard the Holy Grail as incredibly precious; it is therefore tempting to imagine it as being made of gold and jewelry.

The framework presented in this article may serve as basis for a variety of strategies to generate interesting stories. Introducing unexpectedness in an action-to-outcome sequence can indeed be done in different ways. A first idea consists in making both the audience and the character aware of some non-controllable elements in the sequence. The author can decide to introduce a lottery in the sequence. For example, Indiana Jones and the audience are both aware about the risk that the earthenware jug would not be the true Holy Grail. This lottery has to be chosen carefully: the character's action has to be perceived as intentional, and therefore the outcome should remain reasonably expected. Another way to introduce causal unexpectedness consists in introducing intermediary situations between the action and the outcome. By increasing the length of the causal chain, one increases the unexpectedness of the outcome (see (Saillenfest and Dessalles, 2012) for an empirical investigation of how the length of the causal chain may raise interest in stories). A last way to introduce causal unexpectedness is to thwart characters' plans.

These potential developments would present the advantage of being guided by theoretical principles: the global framework of Simplicity Theory, and more precisely the notions of unexpectedness and of intentionality. Their realization will be an opportunity to learn new efficient ways of designing narratives in a (semi-)automatic way, and also to learn more about natural narrative skills.

Funding

This work was supported by the programme *Futur et Ruptures* (Institut Mines-Telecom) and by the *Chaire Modélisation des Imaginaires, Innovation et Création* (<http://imaginaires.telecom-paristech.fr/>).

References

Bae, B. C., and Young, R. M. (2008). *A Use of Flashback and Foreshadowing for Surprise Arousal in Narrative Using a Plan-Based Approach*. In *Proceedings of the 1st Joint International Conference on Interactive Digital Storytelling: Interactive Storytelling, ICIDS 2008*, Erfurt, Germany: 156-67.

Bartsch, K. and Estes, D. (1997). Children's and adults' everyday talk about surprise, *British journal of developmental psychology*, 15(4): 461-75.

Bates, J. (1994). The role of emotion in believable agents. *Communications of the ACM*, 37(7): 122–25.

Bratman, M. E. (1987). *Intentions, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press.

Chater, N. (1999). The Search for Simplicity: A Fundamental Cognitive Principle?, *The Quarterly Journal of Experimental Psychology*, 52A: 273-302.

Chater, N. and Vitányi, P. M. B. (2003). Simplicity: a unifying principle in cognitive science?, *TRENDS in Cognitive Sciences*, 7, 19-22.

Cheong, Y.-G. and Young, R. M. (2008) *Narrative Generation for Suspense: Modeling and Evaluation*, *Proceedings of the 1st Joint International Conference on Interactive Digital Storytelling: Interactive Storytelling, ICIDS 2008*, Erfurt, Germany: 144-55.

Cohen, J. (2001). Defining identification: A theoretical look at the identification of audiences with media characters. *Mass Communication & Society*, 4(3): 245-64.

Dessalles, J-L. (2006). *A structural model of intuitive probability*, *Proceedings of the seventh International Conference on Cognitive Modeling*, Trieste, IT: Edizioni Goliardiche: 86-91.

Dessalles, J-L. (2008). *Coincidences and the Encounter Problem: A Formal Account*, *Proceedings of the 30th Annual Conference of the Cognitive Science Society*, Austin, TX, 2134-39.

Fikes, R. and Nilsson, N. (1971). STRIPS: a new approach to the application of theorem proving to problem solving, *Artificial Intelligence*, 2, 189–208.

Forguson, L. (1989). *Common sense*. New York: Routledge.

Galtung, J. and Ruge, M. H. (1965). The Structure of Foreign News, The Presentation of the Congo, Cuba and Cyprus Crises in Four Norwegian Newspapers, *Journal of peace research*, 2: 64-90.

Gerrig, R. J. (1993). *Experiencing narrative worlds: On the psychological activities of reading*. Yale University Press.

Green, M. C., Brock, T. C., and Kaufman, G. F. (2004). Understanding media enjoyment: The role of transportation into narrative worlds. *Communication Theory*, 14(4): 311-327.

Hidi, S., and Renninger, K. A. (2006). The four-phase model of interest development, *Educational psychologist*, 41(2): 111-27.

Kintsch, W. (1980). Learning From Text, Levels of Comprehension, or: Why Anyone Would Read a Story Anyway. *Poetics*, 9: 87-98.

Labov, W. (1972). The transformation of experience in narrative syntax, In *Language in the Inner City*, Philadelphia: Pennsylvania University Press.: 354–96.

Li, M. and Vitányi, P. (2008). *An Introduction to Kolmogorov Complexity and its Applications*, Springer-Verlag.

Lucas, G. and Meyjes, M. (1989). *Indiana Jones and the Last Crusade*.

<http://www.imsdb.com/scripts/Indiana-Jones-and-the-Last-Crusade.html>

Maguire, R., Maguire, P. and Keane, M. T. (2011). Making sense of surprise: an investigation of the factors influencing surprise judgments, *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 37(1): 176-86.

Maguire P., Moser, P., Maguire, R. and Keane, M. (2013). *A Computational Theory of Subjective Probability [Featuring a Proof that the Conjunction Effect is not a Fallacy]*, *Proceedings of the 35th Annual Conference of the Cognitive Science Society*, Austin, TX, 960-65.

Malle, B. F. and Knobe, J. (1997). The Folk Concept of Intentionality, *Journal of Experimental Social Psychology*, 33: 101-21.

Norrick, N. R. (2000). *Conversational Narrative: Storytelling in everyday talk*. John Benjamins Publishing Company.

Ortony, A. and Partridge, D. (1987). *Surprisingness and expectation failure: what's the difference?*, *Proceedings of the 10th international joint conference on Artificial intelligence, Volume 1*, Los Altos, Morgan Kaufman, 106-08.

Polanyi, L. (1979). So what's the point? *Semiotica*, 25(3-4): 207-42.

Reisenzein, R. (2000). Exploring the strength of association between the components of emotion syndromes: The case of surprise, *Cognition & Emotion*, 14(1): 1-38.

Renninger, K. A. and Hidi, S. (2002). Student interest and achievement: Developmental issues raised by a case study. In A. Wigfield & J. S. Eccles (Eds.), *Development of achievement motivation*. New York: Academic, pp. 173-95.

Riedl, M. O. and Young R. M. (2005). *An Objective Character Believability Evaluation Procedure for Multi-Agent Story Generation Systems*, *Proceedings of the 5th International Working Conference on Intelligent Virtual Agents*, Springer Berlin Heidelberg, 278-91.

Riedl, M. and Young, R. M. (2010). Narrative Planning: Balancing Plot and Character, *Journal of Artificial Intelligence Research*, 39: 217-68.

Saillenfest, A. and Dessalles, J-L. (2012). *Role of Kolmogorov Complexity on Interest in Moral Dilemma Stories*, *Proceedings of the 34th Annual Conference of the Cognitive Science Society*, Austin, TX, August 2012, 947-52.

Saillenfest, A. and Dessalles, J.-L. (2013). Using Unexpected Simplicity to Control Moral Judgments and Interest, In Narratives. In Finlayson, M. A.; Fisseni, B.; Löwe, B. & Meister, J. C. (Eds.), *2013 Workshop on Computational Models of Narrative, Schloss Dagstuhl--Leibniz-Zentrum fuer Informatik*, 32, pp. 214-27.

Schank R. C. (1979). Interestingness: Controlling inferences, *Artificial Intelligence*, 12(3): 273-97.

Schmidhuber, J. (2009). Simple Algorithmic Theory of Subjective Beauty, Novelty, Surprise, Interestingness, Attention, Curiosity, Creativity, Art, Science, Music, Jokes, *Journal of SICE*, 48(1): 21-32.

Schraw G. and Lehman S. (2001). Situational interest: A review of the literature and directions for future research, *Educational Psychology Review*, 13(1): 23-52.

Shaver, K. G. (1985). *The attribution of blame*. New York: Springer-Verlag.

Tannen, D. (1984). *Conversational Style - Analyzing Talk Among Friends*. Norwood: Ablex Publishing Corporation.

Teigen, K. H. and Keren, G. (2003). Surprises: low probabilities or high contrasts?, *Cognition*, 87(2): 55-71.

van Dijk. T. A. (1993). Stories and racism, In *Narrative and Social Control: Critical Perspectives*, Sage Publications, Newbury Park, CA, pp.121-42.

Zwaan, R. A. and Radvansky, G. A. (1998). Situation models in language comprehension and memory. *Psychological bulletin*, 123(2), 162-85.