What killed the cat again? Towards a logical formalization of curiosity (and suspense, and surprise) in narratives with an experimental use of LLM for capturing causality

Benjamin Callac[1], Anne-Gwenn Bosser[1], Florence Dupin de Saint-Cyr[2], Eric Maisel[1]

[1] Lab-STICC_COMMEDIA, CNRS UMR6285, Brest, France
[2] IRIT, Université de Toulouse, France

Résumé

Cet article étend légèrement celui présenté à la conférence TIME'2024[17] dans lequel nous formalisons les affects de la tension narrative – curiosité, suspense et surprise – en logique propositionnelle. Le raisonnement non-monotone est utilisé pour représenter de façon compacte le comportement par défaut du monde et raisonner sur l'état affectif d'un agent. Nous ajoutons ici une expérimentation de calcul automatique de la causalité avec un LLM.

Mots-clés

Représentation de connaissances, Narration, Cognition.

Abstract

This paper slightly extends the one presented at TIME'2024[17] in which we formalize the affects of narrative tension – curiosity, suspense and surprise – in propositional logic. Non-monotonic reasoning is used to compactly represent the default behavior of the world and to reason about the affective state of an agent. We add here an experiment in automatic causality computation with an LLM.

Keywords: *Knowledge Representation, Narration, Cognition*

1 Introduction

Humans tell stories to make sense of the world and communicate their understanding of what happens. Storytelling supposes to be able to sort out which events are worth telling, deciding on a level of detail for describing events, selecting among possible causes the ones which are deemed worth telling. It also supposes to make use of an affective machinery for capturing an audience's attention (emotional contagion, suspense elicitation...). In the act of storytelling, structural and affective phenomena are thus combined with communicative goals in mind. This combination has indeed shown its effectiveness in this respect: the phenomenon of narrative transportation (the experience of being immersed in a story) has been linked to persuasion [27]. The narrative paradigm therefore provides an appropriate framework, in which causal reasoning about the situations narrated [54] is combined with narrative devices to encourage the audience's emotional involvement [51], to study and model how opinion is formed and evolves. Building a framework for reasoning about and unveiling storytelling mechanics could pave the way for intellectual self-defense supporting tools, enabling citizens to arm themselves against hostile disinformation or influence campaigns.

Previous works in structural narratology have studied the way stories are conveyed to their audience and seminal work from (for instance) Genette [25] or Propp [46] have previously served as the backbone inspiration for computational narrative models and storytelling systems [44]. Whilst the operationalization of narrative theories is still subject to debate and caution, such works have shed light on how the story material to tell and the manner in which it is told interacts with a model of the listener[1] (which, depending on the media used for conveying the story can also be a reader, a spectator, or even a gamer): the act of storytelling can thus be understood as knowledge transfer and manipulation of her beliefs.

According to Sternberg [51] or Baroni [5], emotions more specific to narratives which are *suspense*, *curiosity* and *surprise* are critical to retain the interest of the listener. Drivers of the *narrative tension*, they are paramount in maximizing her engagement. In this paper we focus on these narrative tension's building blocks.

In the field of computational narratives, numerous studies and frameworks exist to tell interactive stories, a number of them as an application of planning technologies [12] allowing to adapt the narrative to each user's actions. However, adapting a narrative to a model of the user's emotions remains largely a challenge that needs to be addressed to favor engagement: narrative engagement depends partly on the appropriate maintenance of narrative tension, itself based on the uncertainty occurring in a narrative [8], and listener's models based on a formalization of related emotions have comparatively been less addressed so far in the literature. While suspense and surprise have been the object of previous studies [14] [23], there is — to our knowledge — still no curiosity model applicable to narratives.

In the following, we present a preliminary study

[1] For sake of homogeneity, we use the term listener in all the paper, while this kind of agent is called interpret by Baroni.

for characterizing these emotions from an epistemic standpoint, with a focus on modeling the listener's curiosity depending on her beliefs and knowledge using a propositional language. Our overarching aim is to provide a unifying framework allowing to represent emotions relevant to the characterization of narrative tension and its evolution, which would enable to discuss their relationships and ultimately help establish dramatic metrics about a narrative.

In Section 2, we describe the main emotions supporting narrative tension. We also describe the problems and solutions for formalizing reasoning about action and change, as well as the ways in which the notions of surprise and awareness have been treated in the literature. Section 3 details our proposal, which relies on a non-monotonic framework resulting from extending propositional logic with default rules. The properties of the framework are presented in Section 4, along with some preliminary ideas for developing metrics.

2 Background on reasoning about change and narrative tension

In order to reason about a story, it is useful to dispose of a way to handle the concepts involved for understanding a sequence of facts and events. We first briefly outline the background to this vast subject, before discussing the formalization of emotions in the literature.

2.1 Reasoning about action and change

The formalization of action and change is an old field of research in the domain of knowledge representation and reasoning in AI. There are many different reasoning tasks in this field (see e.g. [18]) like prediction of the new state of the world after an action (which is related to *belief update* [56, 31]), or integration of an observation (which is related to *belief revision* [2]), or *event abduction* which consists in guessing which event took place, or *scenario extrapolation* [19, 16] which consists in taking a partial description of facts and events that occurred and complete it (by prediction or event abduction) or *scenario recognition* [15].

These reasoning tasks were studied in various frameworks, the representation of actions in a compact way has given rise to some problems known as the frame, the ramification and the qualification problems [39, 24, 38]. In propositional logic, these problems were solved by a majority of approaches by introducing a special symbol for expressing a causal rule relating preconditions of an action to its effect (indeed classical implication cannot separate a cause from a consequence due to contraposition). Actions are first described by such rules, then given the set of causal rules, a set of formulas (called frame axioms) are generated stating that any fluent f is true at time $t+1$ if and only if it was (a) true at t and no causal rule concluding $\neg f$ can be fired at t or (b) false at t and a causal rule concluding f can be fired.

As a proof of concept, we choose to use *propositional logic* in this article where we face a problem that can be viewed as an extension of belief extrapolation with narrative tension analysis. Moreover in order to perform non-monotonic reasoning (which allows agents to change their minds and thus accept surprises), we propose to use default rules of the form $a \rightsquigarrow b$ to encode causal relations. Note that the field of non-monotonic reasoning has been extensively studied and many powerful approaches were proposed (see e.g. [29] for an overview), here, we choose to rely on a simple formalism at first.

Logical approaches to computational narratives have been proposed in the past. In [9], (Intuitionistic) Linear logic has been argued to be a suitable representational model for narratives for its capacity to finely represent narrative actions through the production and consumption of resources. This language provides the symbol $\multimap$ which can be used in $A \multimap B$ to express the validity of transforming resource A into resource B, the flow of resources consumption through the associated sequent calculus allowing to establish causal relations. Dynamic logic [30] and its epistemic extensions [7] are formalisms with higher expressiveness. In this work, we propose to characterize narrative tension phenomenon in *propositional logic* (extended with default rules) to demonstrate the representational uniformity of these concepts and their relationships with each other. We will explore how to encode them in aforementioned logics in the future, keeping in mind the challenges raised by their operationalization in their most expressive fragments [3].

2.2 Emotions supporting narrative tension

Psychological models of emotions are often used in the field of affective computing such as models from Ekman [22] or Plutchik [45] (which include surprise). Other works [42] consider that every emotion should have a valence and, as a consequence, surprise, which is inherently neither good nor bad, is considered as a different affective phenomenon. As we position ourselves in the context of studying the emotional states of a listener, we will rely on the characterizations given by Baroni in [5]. Curiosity occurs when there is a partial omission of crucial knowledge: at a given stage when experiencing the storytelling experience, the listener knows they are missing important information. Suspense arises when an event could potentially lead to an impacting result — be it good or bad — to the storyline, and is correlated with anticipation. Surprise results from a rupture from previous expectations, which retroactively invalidates some of the predictions made by the listener: the listener has expectations about how the story will develop, based on story genre or common sense. Going against these expectations while maintaining coherence is what causes surprise. Baroni distinguishes curiosity and suspense from surprise, as the former two are tied to anticipation and an urge to know, whilst the latter arises sporadically as the narrative progresses, which will be reflected in our model.

Related to suspense, the concept of *narrative closure* also reflects the epistemic nature of storytelling (as theorized by Carroll [13]): this encompasses the phenomenological feeling of finality that is generated when all the questions saliently posed by the narrative to the listener are answered.

Previous work in the psychology of narrative understanding [54] has also tied the perception of the importance of story events to causal relationships' perception. In this paper we borrowed from this work, especially tracing a graph representing the narrative with nodes being actions, preconditions and effects and edges being causal relations. We will consider the degree of a node as reflecting its importance in the narrative, reading it as, the more an action is a consequence and has consequences the more important it is.

2.3 Logical models of emotions, surprises and awareness

The logical representation of emotions has already received some attention, see e.g. Lorini [35] or Adam [1] who formalized emotions based on the OCC theory [42]. In these works (which relies on a modal logic for BDI (Belief-Desire-Intention) agents), an agent has beliefs, including beliefs about what is *good for herself*, and expresses different emotions such as joy or sadness.

The particular case of *surprise* was studied by several authors in computer science, but the first study is due to an economist named Shackle [48] who defined the degree of surprise associated with an event as the degree of impossibility of this event given the uncertain knowledge about the situation considered. In Lorini and Castelfranchi [34], the role of surprise is investigated in the context of belief update. They associate a surprise with a difficulty to integrate the new piece of information, this occurs when there is a form of inconsistency between expectation and perception. Surprise was recently formalized in the context of the analysis of jokes by [20], indeed surprise has been considered as an important ingredient for laughter by many authors, the model of surprise of [20] is based on a revision operator and non-monotonic reasoning: to be surprised the listener of a joke should be able to jump to conclusions that can be questioned and even revised.

The characterization of *curiosity* provided by Baroni emphasizes that the listener is aware of its incomplete knowledge and that surprise is linked to a notion of disturbance which makes the agent to question his assumptions/beliefs and leads him to reconsider his understanding of the story. This reconsideration reminds the operation of *awareness raising* introduced by [55] to allow agents "to make their implicit knowledge explicit". Logical models taking into account agent's awareness have previously been defined in the literature. As Halpern [28] states, traditionally when reasoning about agents' beliefs, it is assumed they are aware of every proposition. Modica and Ristichini [40] first came up with a definition of awareness based on knowledge, stating that an agent is aware of p if he knew p or if he knew he did not know p. Halpern extends on this by introducing *implicit* knowledge, where agents are aware of all propositions and can reason with them ; and *explicit* knowledge, which captures the conclusions of which the agent is explicitly aware of. In this system, explicit knowledge is also implicit, while the reverse is not necessarily true.

Previous works have proposed models for agents in computational narratives such as [41] or [11] based on BDI. In [47], a BDI agent aiming to simulate player behavior in interactive stories takes into account the player personality. Other work has assigned personality stereotypes to users [53, 4] according to their interactions with the system. Whilst such models allow personalizing an interactive narrative, they would not enable a storytelling engine to finely drive the narrative tension. By contrast, the *Suspenser* system by [14] offers an operationalization for suspense elicitation, one of the three drivers of narrative tension. In *Suspenser*, suspense is maximized by ordering multiple story bits at the discourse level. We lay in this paper the groundwork for ultimately representing in a unified logical framework suspense, curiosity and surprise, the three drivers of narrative tension. This will build strong foundations for future generative and interactive systems able to operate both at the story and discourse levels.

We approach the modelization of curiosity, suspense and surprise as constructs at given moments of a narrative experience from the listener's beliefs and by non-monotonic reasoning about these beliefs.

Doing so, we believe our model is compatible with previous formalization while providing new insights.

3 Formalizing curiosity, surprise and suspense

We first present an example to illustrate the concepts introduced throughout this article.

Example 1 (The box)**.** *To illustrate the framework, we present a short story involving three agents, Albert, Erwin as well as a protagonist Cecilia*[2] *(respectively agents A, E and C). A short narrative : "Cecilia enters her office. She sees a box lying on her desk that was not there when she last left the room." Our hypothesis is that this event sparks curiosity in Cecilia's mind. We look at it from the point of view of Cecilia who reasons in a closed world where nothing, except three particular events (Albert putting a box on Cecilia's desk, Erwin doing it, Cecilia opening the box) can interact with the state of the world.*

We consider a set of variable symbols $\mathcal{V}$ denoted by Latin lower case letters, from this set of symbols we build the vocabulary $\mathcal{V}_T$ containing all variables of $\mathcal{V}$ indexed by all the integers taken in the set $T = \{0, 1, .., N\}$ representing time points. $\mathcal{L}$ is the propositional language based on $\mathcal{V}_T$ with the usual connectors and constants $\vee$, $\wedge$, $\neg$, $\rightarrow$, $\equiv$, $\perp$ and $\top$ denoting respectively the logical connectors "or", "and", "not", material implication and logical equivalence, contradiction, and tautology. The symbol $\models$ represent satisfiability. Let Ω denote the set of interpretations induced by $\mathcal{V}_T$, we will often use ω for naming a particular interpretation in Ω, each interpretation will be described by the list of literals satisfied by it, e.g., considering the vocabulary $\mathcal{V} = \{a, b\}$, and a set of two time points

[2]We consider a story involving Albert Einstein, Erwin Schrödinger and Cecilia Payne-Gaposchkin, hence the cat in the title.

$T = \{0, 1\}$ $\omega = (a_0, \neg b_0, \neg a_1, \neg b_1)$ is an interpretation in Ω that associates the truth value True to a and False to b at time step 0 and False to a and b at time step 1. The set $Mod(A) \subseteq \Omega$ is the set of interpretations satisfying the set of propositional formulas $A \subseteq \mathcal{L}$ ($Mod(A) = \{\omega \in \Omega \mid \omega \models \bigwedge_{\varphi \in A} \varphi\}$), the same notation is used to represent the set of models of a formula $Mod(\varphi) = \{\omega \in \Omega \mid \omega \models \varphi\}$.

Example 1 (continued): *To study this flow of events taking place in 4 time steps denoted $T = \{0, 1, 2, 3\}$ we need a vocabulary $\mathcal{V} = \{box, A, E, C, empty, vis\}$ meaning respectively there is a box on Cecilia's desk, agent A puts a closed box on the desk, agent E puts a closed box on the desk, agent C opens the box, there is nothing in the box and something inside the box is uncovered (and thus the box has been opened). In the language $\mathcal{L}$ built on $\mathcal{V}$ and T, the following expression is an example of a well-formed formula: $(A_0 \vee E_0) \wedge box_1 \wedge C_2 \wedge \neg empty_2$.*

Default rules are rules that tolerate exceptions and allow us to reason in presence of incomplete information, by assuming that the situation is not exceptional when there is no evidence for the contrary. The notation $\alpha \rightsquigarrow \beta$ (with $\alpha, \beta \in \mathcal{L}$) is used to represent a default rule interpreted as when α is true, it is more plausible that β is true than false.

Example 1 (continued): *In order to be able to encode this example we propose to use default rules to express that by default some fluents keep their value: the following rule is expressing that when there is no box at time point 0 then by default there is no box at time point 1: $\neg box_0 \rightsquigarrow \neg box_1$. This rule admits exceptions: namely, if A puts a box on the desk at time point 0 then generally there is a box at time point 1: $A_0 \wedge \neg box_0 \rightsquigarrow box_1$.*

Given a set of default rules Δ it is possible to define a ranking of these rules according to their specificity, thanks to "System Z" algorithm [43], the default base is then called stratified, its stratas are the formulas with the same rank[3]. Note that there are sets of default rules that do not admit a Z ordering, such default sets are called "inconsistent" in [26]. In this paper, we restrict ourselves to consistent default sets. From a stratified default base lexicographic-entailment [6, 33] is a non-monotonic inference relation which imposes that the more specific the rules, the more mandatory it is to comply with them:

Definition 1 (Lex-inference [6]). *Let $\Delta = \Delta_1 \cup \dots \cup \Delta_n$ be a stratified default base with n strata ordered from the most specific strata Δ_1 to the least specific one Δ_n, and let A and B be two subsets of Δ, and α, β be two formulas of $\mathcal{L}$,*

- Notations: str (for "strict") is a function that translates a set of default rules into a set of formulas of $\mathcal{L}$, i.e., $str(A) = \bigcup_{\alpha \rightsquigarrow \beta \in A} \{\neg\alpha \vee \beta\}$. For all $i \in [1, n]$, and any $E \subseteq \Delta$, E_i denotes the ith strata of E: $E_i = E \cap \Delta_i$.

- A is Lex-preferred *to B given Δ, denoted $A \succ_\Delta B$,*

$$\text{iff there exists } k \in [1, n] \text{ s.t. } \begin{cases} |A_k| > |B_k| \text{ and} \\ \forall i < k,\ |A_i| = |B_i| \end{cases}$$

- A is a Lex-preferred *α-consistent subbase of Δ if $A \subseteq \Delta$ and $str(A) \cup \{\alpha\} \not\models \bot$ and for any $B \subseteq \Delta$ s.t. $str(B) \cup \{\alpha\} \not\models \bot$, $B \not\succ_\Delta A$ holds*

- $\alpha \mathrel{|\!\!\sim}_\Delta \beta$ iff for any Lex-preferred α-consistent subbase B of Δ, $str(B) \cup \{\alpha\} \models \beta$

Example 1 (continued): *Let us consider that the common knowledge Δ about the world consists only in the default persistence of the fluents box, empty and vis and on the default effects of the occurrences of events A, E and C when their preconditions hold.*

$\neg box_0 \rightsquigarrow \neg box_1$	$(A_0 \vee E_0) \wedge \neg box_0 \rightsquigarrow box_1$
$box_0 \rightsquigarrow box_1$	$C_0 \wedge \neg vis_0 \rightsquigarrow vis_1$
$\neg empty_0 \rightsquigarrow \neg empty_1$	$C_0 \wedge \neg vis_0 \wedge empty_0 \rightsquigarrow \neg vis_1$
$empty_0 \rightsquigarrow empty_1$	$\neg box_1 \rightsquigarrow \neg box_2$
$\neg vis_0 \rightsquigarrow \neg vis_1$	...
$vis_0 \rightsquigarrow vis_1$	$C_2 \wedge \neg vis_2 \wedge empty_2 \rightsquigarrow \neg vis_3$

System Z will give a stratification in three strata where all persistence rules (of the form $v_t \rightsquigarrow v_{t+1}$ or $\neg v_t \rightsquigarrow \neg v_{t+1}$) are in the least specific stratum Δ_3 (since they are tolerated by all the other rules). As seen before, $(A_0 \vee E_0) \wedge \neg box_0 \rightsquigarrow box_1$ describes an exception to the persistence of $\neg box$, just as $C_0 \wedge \neg vis_0 \rightsquigarrow vis_1$ describes an exception to the persistence of $\neg vis$ which leads us to place them in Δ_2, the latter itself admits an exception described by rule $C_0 \wedge \neg vis_0 \wedge empty_0 \rightsquigarrow \neg vis_1$ making it the most specific rule thus placed in Δ_1 by System Z algorithm. At the end, we get:

$$\Delta_1 = \{C_t \wedge \neg vis_t \wedge empty_t \rightsquigarrow \neg vis_{t+1}\}_{t \in \{0,1,2\}}$$

$$\Delta_2 = \left.\begin{cases} C_t \wedge \neg vis_t \rightsquigarrow vis_{t+1} \\ (A_t \vee E_t) \wedge \neg box_t \rightsquigarrow box_{t+1} \end{cases}\right\}_{t \in \{0,1,2\}}$$

$$\Delta_3 = \left.\begin{cases} v_t \rightsquigarrow v_{t+1} \\ \neg v_t \rightsquigarrow \neg v_{t+1} \end{cases}\right\}_{\substack{t \in \{0,1,2\} \\ v \in \{box, empty, vis\}}}$$

Using lexicographic inference we get: $\neg box_0 \mathrel{|\!\!\sim}_\Delta \neg box_1$ and $\neg box_0 \wedge (A_0 \vee E_0) \mathrel{|\!\!\sim}_\Delta box_1$, meaning that a priori if there was no box at time 0, there is no box at time 1, but knowing that either A or E has placed a box makes it more plausible that there is a box at time 1.

Note that in this example, for the sake of simplicity, we want to make a closed world assumption (CWA) in order to express that the only possible way to change the variable box (respectively vis) from false to true is the occurrence of A or E (respectively the performance of action C):

$$\begin{aligned} \text{CWA} = \{&(\neg box_t \wedge box_{t+1}) \rightarrow (A_t \vee E_t), \\ &(\neg vis_t \wedge vis_{t+1}) \rightarrow C_t\}_{t \in \{0,1,2\}} \end{aligned}$$

From the set of default rules and the closed world assumption, we can then obtain: $\{box_1\} \cup \text{CWA} \mathrel{|\!\!\sim}_\Delta box_0$ meaning that the most plausible interpretation is that when there is a box at time point 1 it means that there was already a box at time 0. However, if we know that there were no box at time 0 then $\{box_1, \neg box_0\} \cup \text{CWA} \mathrel{|\!\!\sim}_\Delta (A_0 \vee E_0)$

[3]System Z ordering method is based on the tolerance notion between rules. More precisely, a rule $r = \alpha \rightsquigarrow \beta$ is tolerated by a set of n rules $R \subseteq \Delta$ iff $\alpha \wedge \beta \wedge \bigwedge_{\alpha_i \rightsquigarrow \beta_i \in R} (\neg\alpha_i \vee \beta_i)$ is consistent. The process continues until Δ contains only rules tolerated by all the other ones, they constitute the most specific stratum called Δ_1 (Δ_n being the least specific stratum, with n being the number of iterations).

We choose to use the lexicographic entailment in this paper, because [6] have shown that it is a powerful non-monotonic inference relation that satisfies the set of rational properties called System P. The System P, introduced by Kraus, Lehmann and Magidor [32], gathers properties that should follow rationally when one wants to deduce new inferences from a set of existing inference rules. The following definition describes an agent epistemic states via the pieces of information that she believes.

Definition 2 (Agent epistemic state and inference)**.** *A user is represented by a tuple $B = (F, B_{\mathcal{L}}, B_{\Delta})$ composed of a set $F \subseteq \mathcal{L}$ of formulas representing facts, and two sets $B_{\mathcal{L}} \subseteq \mathcal{L}$ and B_{Δ} respectively representing the strict and default rules known by the agent, the default rules of B_{Δ} are expressions of the form $\alpha \rightsquigarrow \beta$ with $\alpha, \beta \in \mathcal{L}$.*
When $F \cup B_{\mathcal{L}}$ and B_{Δ} are both consistent[4], the user is equipped with an inference relation between formulas of $\mathcal{L}$ denoted $\mid\!\sim_B$ defined by:

$$\alpha \mid\!\sim_B \beta \quad \textit{iff} \quad \left| \begin{array}{l} \{\alpha\} \cup F \cup B_{\mathcal{L}} \textit{ is consistent and} \\ \textit{for any Lex-preferred} \\ (\alpha \wedge \bigwedge_{\varphi \in F \cup B_{\mathcal{L}}} \varphi)\textit{-consistent} \\ \textit{subbase} A \in B_{\Delta}, \\ \qquad A \cup \{\alpha\} \cup F \cup B_{\mathcal{L}} \models \beta \end{array} \right.$$

In the following, $\mid\!\sim_B \varphi$ is a shortcut for $\top \mid\!\sim_B \varphi$.

In order to formally introduce curiosity, we need to define awareness. This will be done by simply stating that an agent is aware of a variable if this variable appears in the facts contained in its epistemic state, and we assume that when an agent is aware of a variable then it becomes also aware of every variable of the strict or default rules of its epistemic state containing this variable (mimicking a kind of introspection). We use the notation $v \Subset \varphi$ to express that the variable v appears in the formula φ, this notation can be applied to variables of $\mathcal{V}_T$ as well as of $\mathcal{V}$.

Definition 3 (awareness)**.** *An agent represented by $B = (F, B_{\mathcal{L}}, B_{\Delta})$ is*

- aware of a variable $v \in \mathcal{V}$ *if*
  - *there is a formula $\varphi \in F$ s.t. $v \Subset \varphi$ or*
  - *there is a formula $\varphi \in B_{\mathcal{L}} \cup str(B_{\Delta})$ s.t. $v \Subset \varphi$ and there is a variable $v' \Subset \varphi$ of which the agent is aware; and*
- aware of a formula $\varphi \in \mathcal{L}$ *iff for any variable $v_t \Subset \varphi$, the agent is aware of v.*

Example 1 (continued): *Let us consider that the epistemic state of agent C is $(\emptyset, \mathrm{CWA}, \Delta = \Delta_1 \cup \Delta_2 \cup \Delta_3)$, in this case it does not know any fact, which means that it is not aware of anything. Assume now that at time point 1, our agent Cecilia comes to her office and sees a box on her desk, then the epistemic state of agent C is $(\{box_1\}, \mathrm{CWA}, \Delta)$. In this state, it is aware that a box is on the desk, moreover by introspection its is aware of the possibility to open it due to rule concerning C, the possibility that Albert or Erwin are able to put it on the desk, the possibility that this box is $empty$ or that something inside of it could be vis.*

The following definition enables us to keep only formulas that do not concern a time point later than a given time point t, i.e., keep the formulas such that all their variables are indexed by time points no later than t.

Definition 4 (epistemic state until t)**.** *Given an epistemic state $(F, B_{\mathcal{L}}, B_{\Delta})$ and a time point $t \in [0, N]$, the* epistemic state until t*, denoted $B_{\to t} = (F_{\to t}, B_{\mathcal{L}\to t}, B_{\Delta\to t})$, is defined by:*
$F_{\to t} = \{\varphi \in F \mid \textit{for all } v_{t'} \Subset \varphi, t' \leq t\}$,
$B_{\mathcal{L}\to t} = \{\varphi \in B_{\mathcal{L}} \mid \textit{for all } v_{t'} \Subset \varphi, t' \leq t\}$,
$B_{\Delta\to t} = \{\delta \in B_{\Delta} \mid \textit{for all } v_{t'} \Subset str(\{\delta\}), t' \leq t\}$.

In the following, we use $[\varphi]_{<t}$ (and respectively $[\varphi]_{\leq t}$, $[\varphi]_{>t}$, $[\varphi]_{\geq t}$ and $[\varphi]_{t}$) to denote that φ is a formula containing only variables indexed by time points earlier than t (resp. earlier than or equal to t, strictly later than t, later than or equal to t, equal to t).

Remark 1. *For any formula $\varphi \in F_{\to t} \cup B_{\mathcal{L}\to t} \cup str(B_{\Delta\to t})$, $[\varphi]_{\leq t}$ holds.*

An agent is curious about a formula at time point t if according to its epistemic state until t it is aware of this formula but it is not able to deduce its truth value at time t.

Definition 5 (curiosity)**.** *An agent with state B is* curious about *$\varphi \in \mathcal{L}$ at $t \in T$ if, according to $B_{\to t}$, it is aware of φ and $\not\mid\!\sim_{B_{\to t}} \varphi$ and $\not\mid\!\sim_{B_{\to t}} \neg\varphi$.*

Example 2. *Coming back to Example 1, the epistemic state of C being $B = (\{box_1\}, \mathrm{CWA}, \Delta)$, its state at 0 is $B_{\to 0} = (\emptyset, \emptyset, \emptyset)$, meaning that at 0 it is aware of nothing, thus according to Definition 5 it is not curious about anything at 0. In the epistemic state B, she first thinks that $\{box_1\} \cup \mathrm{CWA} \mid\!\sim_{\Delta} box_0$. However, she remembers that there was no box on her desk at time 0 before she left her office. Meaning that her epistemic state is now $B' = (\{\neg box_0, box_1\}, \mathrm{CWA}, \Delta)$ which enables her to draw the inference $\{\neg box_0, box_1\} \mid\!\sim_{B'} (A_0 \vee E_0)$, however there is no way of knowing which of Albert or Bernard (or both) dropped off the box. More formally, we can say that Cecilia is curious about the possibility that Albert dropped off the box at time 0 because she is aware of this possibility and $\{\neg box_0, box_1\} \not\mid\!\sim_{B'} A_0$ and $\{\neg box_0, box_1\} \not\mid\!\sim_{B'} \neg A_0$. Now if we consider that Albert told Cecilia that he placed a box on her desk at 0 before she entered her office. In that case, the epistemic state of Cecilia is $B'' = (\{A_0, \neg box_0, box_1\}, \mathrm{CWA}, \Delta)$, there is no more ambiguity as she knows who put it there, hence she is not curious about A_0.*

To define suspense, we propose for this formalization to use Baroni's description of *primary suspense* [5] which relies solely on temporal and belief factors. Baroni also describes other types of suspense involving different

[4] Here consistent is not used with the same meaning: for the propositional formulas it means classical logic consistency, while for the default rules base it means that the base can be stratified.

emotional components (empathy and identification with a protagonist for instance). These components affect suspense by strengthening the intensity of curiosity, and we will leave them for further study at the time being. The following definition expresses that an agent feels suspense about a formula φ when this agent is curious about it at time t, and thinks that it is not impossible for facts or events (below denoted ψ) to come to light and reveal the truth of φ (satisfying curiosity about it at last).

Definition 6 (suspense)**.** *An agent represented by an epistemic state $B = (F, B_{\mathcal{L}}, B_{\Delta})$ feels* suspense about $\varphi \in \mathcal{L}$ at time point t *if*

1. *according to B, the agent is curious about φ at time t*
2. *and there is a formula $\psi \in \mathcal{L}$ such that $[\psi]_{>t}$ and $F_{\to t} \cup B_{\mathcal{L}} \cup \{\psi\}$ consistent*
3. *and there is $t' > t$ s.t. either $\mid\!\sim_{B'} \varphi_{t'}$ or $\mid\!\sim_{B'} \neg\varphi_{t'}$ holds, with $B' = (F \cup \{\psi\}, B_{\mathcal{L}}, B_{\Delta})$.*

Example 3. *In the context of Example 1, assume now that agent C has the following epistemic state $B = (\{\neg box_0, box_1, \neg vis_1\}, \text{CWA}, \Delta)$. Here, at time 1 agent C is aware of the box, she is also aware that it is either empty or not, but has no way at this time to know which is true. Formally, $\not\mid\!\sim_B empty$ and $\not\mid\!\sim_B \neg empty$. Hence she is curious about the variable empty at time point 1. Still according to Definition 3, the agent is also aware of the formulas $(C_2 \wedge \neg vis_2 \rightsquigarrow vis_3)$ and $(C_2 \wedge \neg vis_2 \wedge empty_2 \rightsquigarrow \neg vis_3)$. Meaning she is aware she will know the content of the box once she opens it.*
More precisely, the formula $\psi = C_2 \wedge vis_3$ can be added to the facts of the epistemic state because $\{\neg box_0, box_1, \neg vis_1\} \cup \text{CWA} \cup \{C_2 \wedge vis_3\}$ is consistent. Now, $B' = (\{\neg box_0,\ box_1,\ \neg vis_1,\ C_2,\ vis_3\}, \text{CWA}, \Delta)$ yields $\mid\!\sim_{B'} \neg empty_2$. Hence Cecilia feels suspense at time 1 about the truth value of empty.

In order to formalize surprise, following [20], we propose to exploit our non-monotonic setting that enables agents to imagine several more or less plausible situations, i.e., enables them to incorporate new contradicting information by revising previous conclusions. This is required in order to avoid locking the agent in a state of total incomprehension. Surprise can then be defined by the occurrence of a formula that was unexpected but which is completely plausible.

Definition 7 (surprise)**.** *An agent represented by $B = (F, B_{\mathcal{L}}, B_{\Delta})$ is* surprised at time t about a formula $\varphi \in \mathcal{L}$ *if $\varphi \in F_{\to t}$ and $B_{\to t}$ is consistent (φ occurred and it was not impossible) and $B' = (F_{\to t-1}, B_{\mathcal{L}\to t}, B_{\Delta\to t})$ is such that: $\mid\!\sim_{B'} \neg\varphi$ (φ was unexpected)*

Example 2 (continued): *Cecilia is surprised to find the box at time 1. Indeed, given the epistemic state $B = (\{\neg box_0, box_1\}, \text{CWA}, \Delta)$, before seeing the box at time 1, the persistence of $\neg box_0$ into $\neg box_1$ was the most plausible evolution. More formally, we can check that $box_1 \in F$ and $B' = (\{\neg box_0\}, \text{CWA}_{\to 1}, \Delta_{\to 1})$ is such that $\mid\!\sim_{B'} \neg box_1$, and B is consistent (hence $B_{\to 1}$ as well).*

4 Properties

In this section we show several simple properties relating the three emotions, moreover we establish the computational complexity of their detection.
An agent who knows nothing is aware of nothing.

Proposition 1. *If the epistemic state of an agent has no fact, i.e, $B = (\emptyset, B_{\mathcal{L}}, B_{\Delta})$ then the agent is not aware of any variable.*

Proof. Even if $B_{\mathcal{L}}$ or B_{Δ} are non-empty, there is no awareness since no variable appears in F. □

This kind of agent is not curious nor able to feel suspense since curiosity requires awareness, and suspense requires curiosity.

Corollary 1 (of Proposition 1)**.** *If the epistemic state of an agent has no fact, i.e., $B = (\emptyset, B_{\mathcal{L}}, B_{\Delta})$, then the agent is not curious and does not feel suspense about any formula at any time point.*

The following proposition states that an omniscient agent (i.e., an agent with complete information about the world) is never curious nor able to feel suspense.

Proposition 2. *If the epistemic state $B = (F, B_{\mathcal{L}}, B_{\Delta})$ of an agent admits only one most plausible interpretation in Ω, then for any finite formula, there is no time point where the agent is curious or feels suspense about it.*

Proof. In order to be curious, there should exist at least one variable whose truth value is unknown. Hence there should be at least two interpretations that are equally most plausible. □

Because surprise occurs when the agent expects something and then the opposite happens, it means that it is not curious about it (because the surprise makes it know it).

Proposition 3. *Given an epistemic state B of an agent, if the agent is surprised about φ at time t then the agent is not curious about φ neither at time $t-1$ nor at time t.*

Proof. Let us assume that $B = (F, B_{\mathcal{L}}, B_{\Delta})$ is *surprised at time t about a formula $\varphi \in \mathcal{L}$*, it means that $\varphi \in F_{\to t}$ and $B_{\to t}$ is consistent and $B' = (F_{\to t-1}, B_{\mathcal{L}\to t}, B_{\Delta\to t})$ is such that: $\mid\!\sim_{B'} \neg\varphi$. It means that the agent could infer the truth value of φ at time $t-1$, hence she was not curious at $t-1$. Now since $\varphi \in F$ and $B_{\to t}$ consistent then $\mid\!\sim_{B_{\to t}} \varphi$ hence she is not curious about it at t. □

This proposition shows that surprise and curiosity are antagonists in a given epistemic state, however we can imagine stories where the same event sequence may produce curiosity (e.g. by keeping some information hidden, namely the name of the murderer) when told in a given way and surprise when told differently (e.g. revealing this same information at start). The following proposition shows the complexity class of the decision problems associated to awareness, curiosity, suspense and surprise.

Proposition 4. *Given an epistemic state B, a formula $\varphi \in \mathcal{L}$ and a time point $t \in T$,*

- *Deciding whether B is aware of φ at time point t can be done in linear time.*
- *Deciding whether B is curious or feels suspense or surprise about φ at t is Δ_2^p-complete.*

Proof. In order to check awareness about a variable, it is enough to check membership of this variable to a set of formulas, which is linear in the size of the epistemic state, this process should be repeated for all the variables of a formula to check formula awareness. Concerning curiosity, in addition to a test of awareness, it uses two lexicographic inference tests which have been shown to be in P^{NP} by [21]. Suspense requires a curiosity check and a consistency check of the strict part of the base B, which is a SAT problem hence NP-complete. It then requires several lexicographic inferences in order to find the time point where $\mid\!\sim_{B'} \varphi_{t'}$ or $\mid\!\sim_{B'} \neg\varphi_{t'}$ holds. Surprise requires a consistency check of the default base of B (which is a P^{NP}-complete problem according to [21]) and a lexicographic inference, hence the result. □

The complexity P^{NP} of these decision problems is due to the use of the lexicographic inference in their definition. Note that the upper bound (N) on time steps could relieve the computational complexity as obtained in traditional STRIPS planning [10] where the complexity of certain decision problems drops from PSPACE-complete to NP-complete. Note also that formulation of AI planning in answer set programming gives rise to similar complexity [50].

5 Curiosity and causality

For further characterizing narrative tension, we need to quantify the intensity of the emotions generated in an agent when listening to a story. This section is a first attempt towards this goal.

5.1 Curiosity intensity

We propose a definition of the emotional intensity of curiosity. In the following definition we propose to rely on findings from Trabasso and Sperry [54] as a heuristic in order to evaluate the intensity of the curiosity. We first define the causal graph associated with an epistemic state as the one relating variables of $\mathcal{V}_T$ with the links (called E_B) induced by the default rules (B_Δ) and strict rules ($F \cup B_\mathcal{L}$) of the epistemic state B.

Definition 8 (causal graph)**.** *The* causal graph $\mathcal{G}_B$ *induced by an epistemic state $B = (F, B_\mathcal{L}, B_\Delta)$ is a pair (V_B, E_B):*

- $V_B = \{v_t \in \mathcal{V}_T \mid v_t \in \varphi, \varphi \in F \cup B_\mathcal{L} \cup str(B_\Delta)\}$ *is the set of vertices of $\mathcal{G}_B$*
- $E_B = \{(v_t, v'_{t'}) \in \mathcal{V}_T \times \mathcal{V}_T \mid v_t \in \alpha, v_{t'} \in \beta, \alpha \rightsquigarrow \beta \in B_\Delta\} \cup \{(v_t, v'_{t'}) \in \mathcal{V}_T \times \mathcal{V}_T \mid \{l_t\} \cup F \cup B_\mathcal{L} \models l'_{t'}$ *with* $l_t \in \{v_t, \neg v_t\}, l'_t \in \{v'_{t'}, \neg v'_{t'}\}\}$

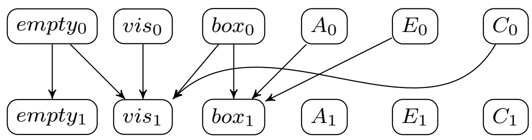

Figure 1: Causal graph induced by the epistemic state $B_{\to 1}$

We illustrate this definition on the epistemic state of Cecilia until time point 1.

Example 3 (continued): *Considering* $B = (\{\neg box_0, box_1, \neg vis_1\}, \mathrm{CWA}, \Delta)$*, the causal graph induced by $B_{\to 1}$ is shown in Figure 1.*

Definition 9 (curiosity intensity)**.** *Given an agent with state $B = (F, B_\mathcal{L}, B_\Delta)$ and* curious about *$\varphi \in \mathcal{L}$ at $t \in T$, her curiosity intensity level is $c_B(\varphi, t) = \sum_{v_{t'} \in \varphi} deg(v_{t'})$ where $deg(x)$ is the degree of the node x in the causal graph induced by B.*

Example 2 (continued): *Given the epistemic state of Cecilia $B = (\{box_1\}$, CWA, $\Delta)$, she is the most curious about vis_1 with intensity 5, denoted $c_B(vis, 1) = 5$. Note that the degree of vis_1 is only four on Figure 1 but there is a supplementary outgoing arc from vis_1 to vis_2 when considering B instead of $B_{\to 1}$.*

Note that another way to refine the curiosity intensity analysis is to take into account the user profile and her preferences over different subjects, which is left for further studies.

5.2 LLMs for building causal graphs

As seen in the previous section the causal graph is crucial to characterize the basic ingredients of narrative tension. In particular we need this graph both for capturing the dynamic laws governing the world and the saliency of each event (hence useful to refine the intensity of curiosity, suspense, and surprise). We first assumed that this graph could be extracted from an encoding of the epistemic state of the agent. Another approach is to take advantage of LLMs in order to obtain this graph, As presented by [52]. While we focused here on a simple story, whose events were transcribed into logic by hand, they tackled the question of generating a narrative graph of salient events from texts. Such system would enable to test our framework on much more complex narratives by automating the generation of such event graphs. Moreover they provide a structured way to identify successively hierarchical, temporal and causal dependencies between events. More precisely, the approach of [52] proceeds in five steps translated in cascading prompts: 1) Summarize the text 2) Capture the events, 3) Generate a hierarchical graph of events, 4) Generate a temporal graph, 5) Generate a causal graph. We[5] experimented their approach on our simple story. It amounted to send a first prompt (corresponding to step 2):

[5] The experiment was conducted by Benjamin Callac a.k.a. "Ben".

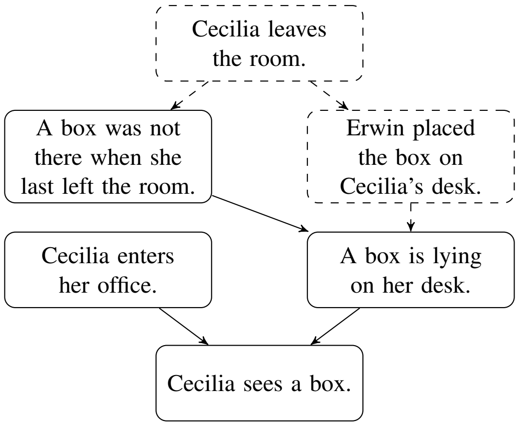

Figure 2: Causal graph generated by an LLM using Tan's [52] method. The two dotted events were added through further prompting of the LLM.

Ben> A structured event is something that happened as described in the text. A structured event is represented as a tuple, which consists of actors, a trigger, and objects. Could you list all the structured events in what follows? "We consider a story involving Albert Einstein, Erwin Schrödinger and Cecilia Payne-Gaposchkin[6] *: Cecilia enters her office. She sees a box lying on her desk that was not there when she last left the room."*

Resulting in the following list of events from ChatGPT-4o[7]:

GPT> Here are the structured events:

1. (Cecilia; enters; her office).

2. (Cecilia; sees; a box).

3. (A box; lying; on her desk).

4. (A box; was not; there when she last left the room).

Following the method used in [52] we then prompt a ChatGPT instance to complete a python program to generate hierarchical, temporal and causal graphs of events. In these prompts the different relations are described as :

Hierarchical Graph: "Each edge in the graph represents a *subevent* relation between the head and tail nodes which are events. They are closely related but on different granularity levels." [8]

Temporal Graph: "Each edge in the graph represents a *temporal* relation between the head and tail nodes which are events. An edge means the head event happens before the tail event temporally."[8]

Causal Graph: "Each edge represents a *causal* relation between the head and tail nodes which are events. An edge means the head event is *caused by* the tail event. The head event will not happen if the tail event did not happen."[8]

We can notice that this causal graph does not contain implicit events (like Albert or Erwin actions), since they were not present in the initial text. Moreover there is no genericity, i.e., the LLM underlies the actual causality in the story but do not give explicitly the causal laws. However each causal arc is labeled with a kind of generic explanation since we ask GPT to add comments for each edge created, e.g. for the edge: `(box_on_desk) -> (cecilia_sees_box)`, it is commented with "Cecilia can only see the box if it is on the desk".

Finally, we observe that the LLM was not initially able to infer events that were not described in the prompt. But, after prompting:

Ben> There are missing events in the story, complete the story with actions from either Albert or Erwin that Cecilia was not aware of.

GPT> (...) Earlier that day, Erwin Schrödinger had placed the box on Cecilia's desk.

Then after asking to complete the graph in python with the event that Erwin put the box there, GPT delivered a python program where it added an edge with the comment that ""A box is lying on her desk." is causally dependent on "Erwin placed the box on Cecilia's desk."". Moreover, we had to suggest to GPT to add the event of Cecilia leaving the room, which finally results in the graph presented in Figure 2.

6 Conclusion

This paper aims at providing a unified framework in which the three emotions at the heart of narrative tension, namely curiosity, surprise, and suspense are formalized and their relationships clarified. This framework is built on non-monotonic reasoning for representing compactly the default behavior of the world and also for simulating the reasoning of an agent in front of a story. The use of non-monotonic reasoning induces a cost in complexity: the detection problems associated with the three emotions are in P^{NP} (due to the use of lexicographic inference). We describe methods to evaluate the curiosity intensity.

While we illustrated our formalization by adopting the point of view of a single agent in a chronological story for the sake of clarity, it does not preclude its adaptability for storytelling using other points of views such as an extradiegetic narrator disclosing knowledge to the listener through a discourse that does not reflect the timeline of the story. To operationalize this model, we plan to investigate different frameworks that are equipped with solvers namely PDDL planning, Linear logic with Ceptre [36] and propositional default logic with TouIST [49]. Moreover, the inherent growing complexity of this problem for scaling to complex narratives requires further study about the granularity of story events, for instance inspired by discussions about the representation of causality [37]. The experiment based on [52] is very promising but further study is needed to establish a generic and guaranteed process for generating valuable causal graphs from any story.

Acknowledgements

We would like to thank the French Defence Innovation Agency (AID) for its support. We also extend our thanks to the reviewers for their insightful comments and ideas for enhancing this paper.

[6]Hence the cat in the title.

[7]"chatgpt-4o-latest" version as of the 15th March 2025.

[8]Prompt extracted from [52].

References

[1] Carole Adam, Andreas Herzig, and Dominique Longin. A logical formalization of the OCC theory of emotions. *Synthese*, 168(2):201–248, May 2009.

[2] Carlos E Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *The journal of symbolic logic*, 50(2):510–530, 1985.

[3] Guillaume Aucher and Thomas Bolander. Undecidability in epistemic planning. In F. Rossi, editor, *Proc. of Int. Joint Conf. on Artificial Intelligence (IJCAI'2013)*, pages 27—33, Beijing, China, 2013.

[4] Heather Barber and Daniel Kudenko. Generation of dilemma-based interactive narratives with a changeable story goal. In *2nd International Conference on INtelligent TEchnologies for interactive enterTAINment*. ICST, 5 2010.

[5] Raphaël Baroni. *La tension narrative: suspense, curiosité et surprise*. Poétique. Éd. du Seuil, Paris, 2007.

[6] Salem Benferhat, Claudette Cayrol, Didier Dubois, Jérôme Lang, and Henri Prade. Inconsistency management and prioritized syntax-based entailment. In *Proc. of Int. Joint. Conf. on Artificial Intelligence (IJCAI'93)*, volume 93, pages 640–645, 1993.

[7] Thomas Bolander and Mikkel Birkegaard Andersen. Epistemic planning for single-and multi-agent systems. *Journal of Applied Non-Classical Logics*, 21(1):9–34, 2011.

[8] Lorenzo Bonoli. Raphaël Baroni, La tension narrative. Suspense, curiosité, surprise, Paris, Seuil, 2007. *Cahiers de Narratologie. Analyse et théorie narratives*, 14, February 2008.

[9] Anne-Gwenn Bosser, Marc Cavazza, and Ronan Champagnat. Linear Logic for Non-Linear Storytelling. *ECAI 2010*, pages 713–718, 2010. Publisher: IOS Press.

[10] Tom Bylander. The computational complexity of propositional strips planning. *Artificial Intelligence*, 69(1-2):165–204, 1994.

[11] Rogelio E. Cardona-Rivera, Bradley A. Cassell, Stephen G. Ware, and R. Michael Young. Indexter : A Computational Model of the Event-Indexing Situation Model for Characterizing Narratives, 2012.

[12] Rogelio E. Cardona-Rivera, Arnav Jhala, Julie Porteous, and R. Michael Young. The story so far on narrative planning. *Proceedings of the International Conference on Automated Planning and Scheduling*, 34(1):489–499, May 2024.

[13] Noël Carroll. Narrative closure. *Philosophical Studies*, 135(1):1–15, August 2007.

[14] Yun-Gyung Cheong and R. Michael Young. Suspenser: A Story Generation System for Suspense. *IEEE Transactions on Computational Intelligence and AI in Games*, 7(1):39–52, March 2015.

[15] Christophe Dousson and Pierre Le Maigat. Chronicle recognition improvement using temporal focusing and hierarchization. In *IJCAI*, volume 7, pages 324–329. Citeseer, 2007.

[16] Florence Dupin de Saint-Cyr. Scenario Update Applied to Causal Reasoning. In *Principles of Knowledge Representation and Reasoning: Proceedings of the 11th International Conference, KR 2008*, pages 188–197, January 2008.

[17] Florence Dupin de Saint-Cyr, Anne-Gwenn Bosser, Benjamin Callac, and Eric Maisel. What killed the cat? Towards a logical formalization of curiosity (and suspense, and surprise) in narratives. In Michael Sioutis, Pietro Sala, and Fusheng Wang, editors, *LIPIcs*, volume 318, pages 10:1–10:16, Montpellier, France, October 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[18] Florence Dupin de Saint-Cyr, Andreas Herzig, Jérôme Lang, and Pierre Marquis. Reasoning About Action and Change. In Pierre Marquis, Odile Papini, and Henri Prade, editors, *A Guided Tour of Artificial Intelligence Research*, volume 1 / 3 of *Knowledge Representation, Reasoning and Learning*, pages 487–518. Springer International Publishing, May 2020.

[19] Florence Dupin De Saint-Cyr and Jérôme Lang. Belief extrapolation (or how to reason about observations and unpredicted change). *Artificial Intelligence*, 175(2):760–790, February 2011.

[20] Florence Dupin de Saint-Cyr and Henri Prade. Belief revision and incongruity: Is it a joke? *Journal of Applied Non-Classical Logics*, 33(3-4):467–494, October 2023.

[21] Thomas Eiter and Thomas Lukasiewicz. Default reasoning from conditional knowledge bases: Complexity and tractable cases. *Artificial Intelligence*, 124(2):169–241, 2000.

[22] Paul Ekman. An argument for basic emotions. *Cognition and Emotion*, 6(3-4):169–200, May 1992.

[23] Jeffrey Ely, Alexander Frankel, and Emir Kamenica. Suspense and Surprise. *Journal of Political Economy*, 123(1):215–260, February 2015.

[24] Joseph Jeffrey Finger. *Exploiting constraints in design synthesis*. PhD thesis, Stanford University, Stanford, CA, 1987.

[25] Gerard Genette. *Nouveau Discours du Récit*. Seuil, 1983.

[26] Moisés Goldszmidt and Judea Pearl. On the consistency of defeasible databases. *Artificial Intelligence*, 52(2):121–149, 1991.

[27] Melanie Green and Timothy Brock. The Role of Transportation in the Persuasiveness of Public Narrative. *Journal of personality and social psychology*, 79:701–21, November 2000.

[28] Joseph Y. Halpern. Alternative Semantics for Unawareness. *Games and Economic Behavior*, 37(2):321–339, November 2001.

[29] Andreas Herzig and Philippe Besnard. *Knowledge Representation: Modalities, Conditionals, and Nonmonotonic Reasoning*, pages 45–68. Springer International Publishing, Cham, 2020.

[30] Kaarlo Jaakko Juhani Hintikka. *Knowledge and belief: An introduction to the logic of the two notions*. Cornell University Press, Ithaca and London, 1962.

[31] Hirofumi Katsuno and Alberto O Mendelzon. On the difference between updating a knowledge base and revising it. *KR*, 91:387–394, 1991.

[32] Sarit Kraus, Daniel Lehmann, and Menachem Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial intelligence*, 44(1-2):167–207, 1990.

[33] Daniel Lehmann. Another perspective on default reasoning. *Annals of mathematics and artificial intelligence*, 15:61–82, 1995.

[34] Emiliano Lorini and Cristiano Castelfranchi. The cognitive structure of surprise: looking for basic principles. *Topoi*, 26(1):133–149, 2007.

[35] Emiliano Lorini and François Schwarzentruber. A logic for reasoning about counterfactual emotions. *Artificial Intelligence*, 175(3):814–847, March 2011.

[36] Chris Martens. Ceptre: A language for modeling generative interactive systems. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 11, pages 51–57, 2015.

[37] Lawrence J. Mazlack. Granular causality speculations. In *IEEE Annual Meeting of the Fuzzy Information, 2004. Processing NAFIPS '04.*, volume 2, pages 690–695 Vol.2, 2004.

[38] John McCarthy. Epistemological problems of artificial intelligence. In *Proc. Int. Joint Conf on Artificial Intelligence (IJCAI'77)*, pages 1038–1044. Elsevier, 1977.

[39] John McCarthy and Patrick J Hayes. Some philosophical problems from the standpoint of artificial intelligence. *Machine Intelligence*, 4:463–502, 1969.

[40] Salvatore Modica and Aldo Rustichini. Awareness and partitional information structures. *Theory and decision*, 37:107–124, 1994.

[41] Emma Norling. Capturing the quake player: Using a BDI agent to model human behaviour. In *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 1080–1081, Melbourne Australia, July 2003. ACM.

[42] Andrew Ortony, Gerald L Clore, and Allan Collins. *The cognitive structure of emotions*. Cambridge university press, 2022 (original work published 1988).

[43] Judea Pearl. System Z: A natural ordering of defaults with tractable applications to nonmonotonic reasoning. In *Proc. 3rd Conf. on Theoretical Aspects of Reasoning about Knowledge*, pages 121–135, 1990.

[44] David Pizzi, Fred Charles, Jean-Luc Lugrin, and Marc Cavazza. Interactive Storytelling with Literary Feelings, April 2007.

[45] Robert Plutchik. A general psychoevolutionary theory of emotion. In *Theories of emotion*, pages 3–33. Elsevier, 1980.

[46] Vladimir Propp. *Morphology of the Folktale: Second Edition*. University of Texas Press, 1968.

[47] Jessica Rivera-Villicana, Fabio Zambetta, James Harland, and Marsha Berry. Using BDI to Model Players Behaviour in an Interactive Fiction Game. In Frank Nack and Andrew S. Gordon, editors, *Interactive Storytelling*, volume 10045, pages 209–220. Springer International Publishing, Cham, 2016.

[48] George Lennox Sharman Shackle. *Decision, Order and Time in Human Affairs*. (2nd edition), Cambridge University Press, UK, 1961.

[49] Khaled Skander Ben Slimane, Alexis Comte, Olivier Gasquet, Abdelwahab Heba, Olivier Lezaud, Frederic Maris, and Maël Valais. Twist your logic with TouIST. *CoRR*, abs/1507.03663, 2015.

[50] Tran Cao Son, Enrico Pontelli, Marcello Balduccini, and Torsten Schaub. Answer set planning: a survey. *Theory and Practice of Logic Programming*, 23(1):226–298, 2023.

[51] Meir Sternberg. How Narrativity Makes a Difference. *Narrative*, 9(2):115–122, 2001.

[52] Xingwei Tan, Yuxiang Zhou, Gabriele Pergola, and Yulan He. Cascading large language models for salient event graph generation. In Luis Chiruzzo, Alan Ritter, and Lu Wang, editors, *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 2223–2245, Albuquerque, New Mexico, April 2025. Association for Computational Linguistics.

[53] David Thue and Vadim Bulitko. Modelling goal-directed players in digital games. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 2, pages 86–91, 2006.

[54] Tom Trabasso and Linda L Sperry. Causal relatedness and importance of story events. *Journal of Memory and Language*, 24(5):595–611, October 1985.

[55] Johan Van Benthem and Fernando R Velázquez-Quesada. The dynamics of awareness. *Synthese*, 177:5–27, 2010.

[56] Marianne Winslett. *Updating Logical Databases*. Cambridge University Press, 1990.