Basics of Narrative Interpretation: Physical Model and Character-Specific Views of the Storyworld*

Pablo Gervás*

Facultad de Informática, Universidad Complutense de Madrid, Madrid, 28040 Spain

Abstract

In many stories, an important feature that contributes to their interest is the existence of more than one possible interpretation of the facts of the story. Authors often play with these ambiguities by making both interpretations available to the reader, but only one of them to different subsets of the characters. This sometimes gives rise to a conflict between the views held by different characters that is only resolved towards the end of the story. To adequately model computationally the impact of such stories on the reader, we need not only to be able to construct a representation of each of the possible views on the story, but also to keep track of which characters hold which view, when conflicts between them arise, and at which points of the story the conflicts are resolved. The present paper explores one possible computational model that addresses this task, and explores the insights it gives rise to when applied to a certain type of story.

Keywords

model of the reader, stories with more than one interpretation, tracking conflicting views on a story, death of the author, birth of the reader

1. Introduction

When we read a story we do much more than reconstructing the set of events that happened to the protagonists. There is an empathetic side to our nature that helps to automatically build a fair representation of what particular characters are feeling in the face of what happens to them. In fact, our minds are trained to build at the same time representations for the feelings of many different characters that appear in the story [1]. In many cases, our reaction to the story at an emotional level is based more on the contrasts between these representations than either the set of events being reported or even our own reaction to those events. Yet in terms of computational modeling of the way we process narrative, we know very little about the mechanics that are necessarily involved in the task of isolating the views on the story that each character has, storing them as we proceed along the discourse, and comparing them to inform our emotional reactions to the story.

The present paper explores some of these issues at a basic level, proposing a computational model that emulates the process of interpreting an outline of a story first into a model of the world at the physical level [2] and then into a set of distinct views on the story as perceived by the different characters involved [1]. To avoid the risk of getting bogged down in the challenges of parsing natural language text the outline of the story considered as input is represented in terms of conceptual representations of the events that are easy to interpret computationally. The model of the world at the physical level allows us to establish which characters are aware of which events of the story, which is crucial to work out the view of the story that each one perceives. Based on the subset of the events that each character perceives, a different view of the story is constructed based on an existing model for the computational interpretation of narrative [3]. This model includes capabilities for dealing with embedded stories, using these embedded stories as anaphoric references to prior sequences of events, identifying spans of discourse that present alternative views on the same set of events, and allowing some of these embedded stories to be tagged as insincere.

8th international workshop on Computational Models of Narrative, 28-30 May 2025, University of Geneva - Switzerland
*Corresponding author.
pgervas@ucm.es (P. Gervás)
http://nil.fdi.ucm.es/ (P. Gervás)
0000-0003-4906-9837 (P. Gervás)

The resulting model is tested over a simple story that includes a sufficient number of interesting phenomena to showcase the type of insights on a story that may be derived from the proposed approach to processing it.

2. Previous Work

A number of topics need to be reviewed to provide context for the work in the present paper: psychological accounts of how readers process fiction, views of the role of ambiguity in narrative, the importance of character point of view, and existing computational models for the interpretation of narrative.

2.1. The Process of Reading Fiction

Psychologists have proposed some basic descriptions of how people's minds react to the reading of fiction. There is a part of the process that addresses the construction of physical model of the world [2, 4]. This physical model of the world acts as a basis for the construction of representations of the mindstates of the storyteller and the characters in the story [1, 5]. In this second process the reader appears to build separate models of the view of the storyworld that each character in the story has.

2.2. Ambiguous Interpretations in Narrative

Vernant [6] points out that a fundamental structural device in classical Greek tragedy is the existence of more than one possible interpretation of what the characters are saying on stage. In the face of this ambiguity, the hero will be seen to understand only one interpretation, and to cling to it even though it leads him to perdition. The choir, in contrast, usually perceives both intepretations and comments on them as the play progresses. Vernant postulates that the fact that the spectator can perceive both interpretations from the start, and is aware of the fact that the hero has not, is a fundamental aspect of this type of tragedy. Barthes [7] relies on this idea to argue for a raise in the significance of the role of the reader in the processing of a literary text. Barthes states the "the total existence of writing" implies that a text is made of "multiple writings", and that it is the role of the reader to identify them all, work through them, and build her own impression of the text based on that process.

2.3. Character's Point of View

Friedman [8] argues that since the time of the Victorian novel, in which the presence of the author was prevalent – as narrator but often also as commentator and interpreter of the events in the story – fiction has moved on to focus more specifically on the point of view of characters. In this new approach "the reader perceives the action as it filters through the consciousness of one of the actors involved". This enables the author to represent the "prejudices and predispositions" of her characters and to "evaluate [them] dramatically in relation to one another within their own frame".

Shanahan and Shanahan [9] propose Character Perspective Charting as a procedure to help children get more information out of the processing of a story. Instead of restricting the analysis of a story to the point of view of the main character, this procedure encourages the construction of distinct views – perspectives – of the story for other characters involved. Shanahan and Shanahan suggest that this technique is best used with stories that "have two or more characters with separate goals". An appropriate set of characters is selected and a separate chart is built for each one of them. A chart compiles for a given character elements such as the problem that they face in the story, their goals, their intentions and the actions they undertake to achieve them.

2.4. Computational Interpretation of Narrative Discourse

Ricoeur [10] pointed out that there is value in exploring the way we use language for pragmatic purposes beyond the the set of "grammatical devices" that particular languages employ to implement the two basic functions of a sentence: to establish what is being talked about and to establish what is predicated about it. A sentence can be characterised by a predicate that applies to some entities identified by the subject (and objects). A discourse is then defined as a sequence of such sentences that is an act of communication that brings together one actor producing the discourse – a speaker or writer – and (at least one) other actor consuming it – the listener or reader. In this paper we focus on the process of interpreting a sequence of simple predicates that present the outline of a story – which qualifies as a discourse in Ricoeur's sense – into computational data structures that capture some of the features that we know mental representations built by reader may have.

It is at this level of representation that Ryan [11] phrases her argument in favour of the modal structure of narrative universes. Starting from a description of a narrative plot as a temporal succession of different states of affairs mediated by events, in which states and events are represented in terms of sets of propositions considered true at the relevant times, Ryan argues for a model of narrative based on a number of possible worlds. Of these possible worlds, one of them describes the 'actual world' of the narrative – which contains the events considered to be true in the storyworld – and the others describe alternative views of the storyworld which are predicated in a different mode – as facts that certain characters in the world either know, consider as possibilities, desire, intend or view as moral imperatives. Based on these concepts, Ryan proposes that the structure of a narrative may be described as a sequence of movements of the set of worlds in a narrative universe, over which the characters involved attempt to make the actual world match as best as possible their personal model worlds. In her conclusions, Ryan states that the model she has described "suggests" an outline for an "algorithm" for the cognitive processing of narrative discourse. This algorithm traverses the narrative discourse and progressively builds a data structure that combines a series of states and a related series of chronologically ordered events. At each point in the discourse the following data have to be identified: the narrator and the potential hearers of a given sentence, and the set of possible worlds for the narrative universe that should be affected by the sentences – the actual world, the worlds that represent what narrator and hearers know and the worlds that are referred to if the sentence is an event of linguistic communication. Based on those data the sentences is used to update the possible worlds in the set identified.

Gervás [3] had already developed a computational model for the interpretation of discourse that operated on similar simple inputs and captured basic interpretation of *embedded stories*,[1] including their uses as anaphoric devices for referring to spans of discourse already presented to the reader, as means of introducing flashbacks to events earlier in the story but not yet presented to the reader [13], and as means of introducing alternative versions of spans of discourse already told.

3. Computational Tools for Isolating Character-Specific Views

This section presents: a simplified format for conceptual description of narrative discourse to be used as input for the rest of the steps, the basic algorithm employed for interpreting narrative for a single view, the algorithm proposed for constructing a physical model of the world required to identify which characters are present at each point in the narrative, and the algorithm proposed to identify the different views that different characters have of a given story.

3.1. A Simplified Input for Narrative Discourse

The input format employed for story outlines in the rest of the paper follows the definitions in [3]. The basic elements are summarised here to make the paper easier to understand.

- A *discourse* is represented as a sequence of updates to the story.
- An *update* is represented by a set of statements.
- A *statement* is a *predicate* followed by an open number of arguments. Arguments in a predicate are filled with *constants* that identify the elements in the story world to which the predicate applies.

[1] The "literary device of the 'story within the story', the structure by which a character in a narrative text becomes the narrator of a second narrative text framed by the first one" [12].

Discourse	**Interpretation** Narrative level 0	 Narrative level 1
⋮	⋮	
start_story princess'_torment kidnap dragon princess torment_at_night dragon princess		
tell_story X brother2 princess'_torment	tell_story X brother2	**princess'_torment** kidnap dragon princess torment_at_night dragon princess
decides_to_react brother2 ⋮	decides_to_react brother2 ⋮	

Table 1
Example of input discourse sequence and structured interpretation for tale 155 as analysed by [15]. The left hand column the shows how embedded story is explicitly declared in the discourse sequence used as input: the inline declaration of the embedded story, and the statement that introduces the embedded story have been boxed for clarity. The right hand column shows how the embedded story is explicitly separated as a distinct structure within the interpretation.

- The set of statements in an update is understood to represent the conjunction of the corresponding predicates.

Embedded discourses are represented as distinct spans of the input discourse within the larger structure of the discourse, along the lines followed in [14]. An embedded story will be represented as a combination of the following elements: (1) an update that includes a specific statement to act as marker in the discourse for the beginning of an embedded story, with an argument that defines a name for the embedded story to act as its unique id (`start_story <story-name>`), (2) a sequence of updates for that story (just like those for the frame story), (3) an update that includes a statement to convey the telling of the embedded story within the frame story (`tell_story <narrator> <narratee> <story-name>`) which also acts as as marker in the discourse for the end of the embedded story.

Table 1 presents an example of of the discourse sequence used as input – left hand column – and the interpretation that is constructed from it – right hand column.

3.2. Basic Interpretation Algorithm

The processing of a story into a single view is done using the algorithm proposed in [14] for the interpretation of embedded stories.[2]

3.2.1. Segmenting into Narrative Levels

The first task to be addressed is the segmentation of the discourse into the various narrative levels involved in it. This is done by relying on a stack-based mechanism for handling the changing contexts of interpretation, as described in Table 2.

The outcome of this algorithm is a *narrative-level segmented discourse* representation for the narrative that is recursively nested into as many layers as the number of narrative levels present in the story. The stack should be empty at the end. Each narrative level is represented by a different story interpretation: the frame story in the main interpretation, those for embedded stories in the representations stored in the table. The recursive nature of the embedded stories is captured by the presence in the corresponding

[2]To ensure the process described in the rest of the paper is understandable without resorting to the original paper, a brief description is included here. Interested readers are invited to consult the original paper for further details.

- start with empty story interpretation for frame story, empty stack for initial narrative level, and empty table of embedded stories
- on start of an embedded story (`start_story <story-name>`):
  - push to stack interpretation of frame story so far
  - create new empty story interpretation for embedded story
- process updates for embedded story onto story interpretation for embedded story
- on end of embedded story (`tell_story statement <narrator> <narratee> <story-name>`):
  - store accumulated interpretation for embedded story in table for embedded sub-stories indexed by name of sub-story (`<story-name>`)
  - pop from stack interpretation for frame story acting as context, establish it as context for rest of frame story
  - add special `tell_story <narrator> <narratee> <story-name>` statement to interpretation of frame story to encode how telling of embedded story fits into frame story

Table 2
Basic algorithm for interpreting embedded stories.

narrative level of telling predicates that refer in each case to the index of the corresponding embedded story in the table.

The representation to this point partitions the input narrative discourse into its constituent sub-stories, but it does not capture the relations that may connect together the different sub-stories. Some of the sub-stories may simply not be connected at all to the frame story (*unrelated stories*). Some of the sub-stories may be connected to the frame story by referring to events that have happened in the same storyworld as the frame story but which had not been mentioned before (*preceding stories*). Some of the sub-stories may be connected to the frame story by referring to the events in the storyworld for the frame story that have already been mentioned in the preceding discourse (*anaphoric stories*), or by presenting alternative versions of some events in the storyworld that have already been told (*conflicting stories*). Some of the stories may refer to states of the world other than the actual one (*modal stories*). When sub-stories refer to events in the same storyworld, an important relation between them that needs to be established is the relative chronology. Anaphoric stories often appear in a story when events already told in the frame story are retold by a witness to someone who was not present. Preceding stories are the main tool used by authors to introduce flashbacks.

3.2.2. Estimating Relations Between Sub-stories

The narrative-level segmented discourse representation as it stands still fails to capture information that humans intuitively extract during interpretation, such as: (1) the relative order – in storyworld chronology – of embedded stories that occur in the same storyworld as the frame story, (2) instances where an embedded story refers to a span of the narrative already seen by the reader – one character tells another about past exploits – (3) instances where an embedded story overlaps and provides a different view of the events than an existing prior span.

A further step of processing of the *narrative-level segmented discourse* is required to identify:

- a plausible relative ordering between the events mentioned and
- any conflicting view on particular sets of events

The events in the main frame story and the set of sub-stories in the *narrative-level segmented discourse* are processed to produce a *branching partially ordered graph.* The events in the main frame story are compiled into an ordered sequence that acts as reference for the placement of the sub-stories

- insert events from *frame story* into graph after preceding event in frame story
- on reaching embedded story:
  - if embedded story involves characters not present in frame story, mark as *unrelated story* and store separately
  - otherwise: search preceding spans of frame story for matches:
    - if match is found, mark embedded story as *conflicting story* and insert into graph before start of matching span and marked as conflicting view on events in the span
    - otherwise, mark embedded story as *preceding story* and insert into graph before start of frame story (it refers to a time before that point)

Table 3
Algorithm for placing fabulae for embedded stories with respect to fabula from frame story.

with respect to them. The procedure used for insertion of the various elements into the graph and classification of the sub-stories is described in Table 3.

This procedure is recursive, as every embedded story has a potential for being a frame story for further embedded stories. As the same representation for the frame story is used for the embedded stories, the recursion is handled by construction. This algorithm produces a *branching partially ordered graph* that holds all the events in the narrative.

3.3. Building a Physical Model of the Storyworld

The steps described above partition the narrative discourse into a set of distinct narrative levels, and rely on some simple heuristics to identify – to a certain extent – whether the events mentioned at different narrative levels relate to the same storyworld, and, if they do, what their relative temporal order might be.

However, the representation that results for each narrative level is still a linear sequence of statements that does not capture any information about the physical configuration of the world in which the story happens. In contrast, human interpretation of stories immediately yields a certain intuitive understanding of the physical spaces in which the story is happening, at the very least in terms of a set of locations that the characters move between.

An important challenge is that the model of the world to be constructed must necessarily be partial. Much as we would like to be able to reconstruct directly a physical map of the world – equivalent to the one that might be available to an author attempting to create a story –there is a loss of concrete information about the physical details of the world in the process of encoding a storyworld as a narrative. The evidence available to the decoder of a narrative discourse is by nature partial, and this has to be taken into account in the representation of the model that such a decoder may build. Human readers manage to infer from the data present in a narrative discourse a lot of the information that has been removed during encoding. They do this by applying common sense. This is known to be very hard to emulate computationally.

Nevertheless, it is possible to enrich the representation of the input that we are considering with additional information regarding locations that may inform a simple model we want to build of the process of predicting what events each character is aware of. To achieve this:

- additional predicates for indicating spatial complements to the sentences are introduced (`at_location` to indicate an origin location and `to_location` to indicate a destination location)
- the input discourse is rewritten to include such spatial complements in the story updates that involve a change in location

Story updates	WorldSnapshot	LocationFlash 1	LocationFlash 2
start_story *princess'_torment*	1		
at_location palace kidnap dragon princess2 misbehaved dragon **to_location dragons_lair**	2 3	**palace** []	**dragons_lair** [dragon, princess2]
torment_at_night dragon princess2	4	**dragons_lair** [dragon, princess2]	
tells_a_story someone brother2 *princess'_torment* story princess'_torment	5	**loc-1** [someone, brother2]	
decides_to_react brother2	6	**loc-1** [brother2]	
sets_out brother2 **to_location dragons_lair**	7	**loc-1** []	**dragons_lair** [brother2]
fight brother2 dragon	8	**dragons_lair** [brother2,dragon]	
wounded brother2	9	**dragons_lair** [brother2]	
defeats brother2 dragon	10	**dragons_lair** [brother2,dragon]	
releases brother2 princess2	11	**dragons_lair** [brother2,princess]	
sets_out princess2 **to_location way_to_palace**	12	**dragons_lair** []	**way_to_palace** [princess]

Table 4
Example of input narrative discourse with additional predicates describing location-related information. Spatial complement predicates are marked in bold. Columns on the left show alongside the dynamic model built for it.

- some mechanism must be made available to identify predicates that involve a change in location (a simple look up list is sufficient)

An example of discourse enriched in this way is shown in Table 4. The kidnapping of the princess by the dragon involves a transfer from the palace to the dragon's lair, the hero moves to the dragon's lair to confront the dragon, and the princess once released sets out on her way to the palace.

To process such inputs, we want to consider a model in two stages:

- a representation of the hard evidence that the decoder finds in the narrative discourse
- a richer model that is inferred from the hard evidence by applying some simple rules

The richer model would be closer to the type of physical map that an author might have in mind. But it will necessarily be changing as the narrative discourse for the story is parsed, because more information about the world may be gleaned with each further update to the story.

The input can now be traversed to mine this new layer of location-related information into a very basic model for the hard evidence on the physical characteristics of the world that might be obtained at each step from the story updates in the narrative discourse.

The model of the physical world is based on the concept of a `LocationFlash`, a data structure that holds:

- a label identifying a location in the world (as mined from the spatial complement predicates)
- a set of labels for characters mentioned in the story as present in that location

The *dynamic model* of the world (`DynamicWorldModel`) consists of a sequence of `WorldSnapShot` that each holds:

- a label corresponding to a story update from the discourse
- a set of `LocationFlash` (only one for most story updates, two for story updates involving movement across locations)

The following heuristics are applied to process the sequence of updates in an input discourse:

- (if the update does not involve movement) create a new `WorldSnapShot` with a `LocationFlash` for the given location and any characters that are mentioned in the update
- (if it does) create a new `WorldSnapShot` with two `LocationFlash`: one for the start location with the characters that remain in it, and one for the destination location with the characters that have moved
- add the created `WorldSnapShot` to the `DynamicWorldModel`

The result of this process is a poor representation of the physical structure of the storyworld, but it has the significant advantage of including only factual information that has been explicitly made available in the discourse. That means that the information it contains will be consistent over the discourse, and not likely to change as a result of inferences made from later contributions to the discourse.

Table 4 shows an example of the `DynamicWorldModel` constructed for a simple input discourse. The example illustrates how spatial complements can be used to express both origin and destination location, and how in the absence of information to the contrary the preceding location is carried forward. The telling of a story may imply a change in location in as much as the narrative shifts from the narrative level of the embedded story – told by someone to the brother, about the kidnapping of the princess – to the narrative level of the frame story, where the brother learns about it. In the cases where a change of location is identified but the name of the location is not given, a generic identifier for the location is inserted, to allow the system to keep track of it.

Over such a representation, we can apply a further set of heuristics to expand the information into a more detailed representation of the world. However, this representation based on inference is only valid as a tentative model of the physics of the storyworld at the point in the discourse in which it has been inferred. Subsequent updates will modify the `DynamicWorldModel` and may lead to different inferences at later stages.

The heuristics applied to construct a more detailed (inferred) model of the world from the `DynamicWorldModel` are based on the following assumptions:

- locations that exist in snapshots of the world at some point in the discourse still exist in other snapshots for the same world
- characters that were present in a given location in prior snapshots of the world will still be present at the same location in later snapshots unless the discourse has explicitly mentioned that they have moved

These assumptions constitute applications of Grice's maxims of conversation [16], which state that if the author wanted us not to make such assumptions, she would have included in her discourse some evidence to block them.

Based on these assumptions the instances of `WorldSnapShot` in the `DynamicWorldModel` can be expanded so that:

- every `WorldSnapShot` in the `DynamicWorldModel` includes instances of all the locations that have been mentioned to this point in the discourse
- the continuity of all the characters mentioned is enforced (there is no `WorldSnapShot` in which a character that existed before or after has disappeared)
- the consistency of `DynamicWorldModel` is enforced (there is no `WorldSnapShot` in which a character appears at the same time in more than one location)

Table 5 shows an example of the more refined representation of the world inferred from the `DynamicWorldModel` in Table 4. The gaps in continuity that were apparent in the dynamic model have been filled in, and the representation now reflects that movement of characters across the locations in the world: dragon and princess from palace to dragon's lair, brother to dragon's lair, princess away when she is released.

Time points	Locations **palace**	**dragons_ lair**	**loc-1**	**way_to_palace**
1	[dragon, princess2]	[]	[someone, brother2]	[]
2		[dragon, princess2]	[someone, brother2]	[]
3		[dragon, princess2]	[someone, brother2]	[]
4		[dragon, princess2]	[someone, brother2]	[]
5		[dragon, princess2]	[someone, brother2]	[]
6		[dragon, princess2]	[someone, brother2]	[]
7		[dragon, princess2, brother2]	[someone]	[]
8		[dragon, princess2, brother2]	[someone]	[]
9		[dragon, princess2, brother2]	[someone]	[]
10		[dragon, princess2, brother2]	[someone]	[]
11		[dragon, princess2, brother2]	[someone]	[]
12		[dragon, brother2]	[someone]	[princess2]

Table 5
Example of representation of the world inferred from the DynamicModel.

3.4. Identifying Different Views on the Story by Different Characters

The *branching partially ordered graph* described at the end of Section 3.2 captures significant information about the events that happen in the story. But it constitutes a single generic view for all the events, from the perspective of an omniscient narrator. To obtain insights into how different characters might have different views of the story, we would need to be able to isolate the specific subsets of the graph that each character has access to in the course of the narrative.

This corresponds to being able to explore the same story under different focalizations. The concept of *focalization* [17] refers to the way people tell stories by focalizing on a given character at a time, so that the reader learns about the view of the world that this character has, foregoing the information that might have been available to other characters.

What we are proposing is to rely on our enriched model of the physical storyworld to inform at each point in the discourse which characters are in a situation to actually perceive the events being described, and which ones are not. Based on this information we can now build not a single branching partially ordered graph for the whole story, but a set of branching partially ordered graphs, one for each different character.

To explore the challenges that this task presents, we will consider a more elaborate discourse that presents more than one alternative view of the storyworld, and involves different characters shifting in their perception of what is true as the story progresses. Table 7 shows the full transcription of the input discourse for this story, which corresponds to an extract of tale 155 of the corpus analysed by Propp, which includes an instance of his false hero sequence.

Table 7 shows the schematic of the plot lines of the different embedded stories present in tale 155. This is a transcription of the topological structure identified by the algorithm in the branching partially ordered graph. This schema corresponds to the view of the story that a reader, or an omniscient narrator may have. It includes simultaneously all the various views on the story that are presented at some point of the discourse. This is less than useful in at least two different ways.

First, because in the process of reading the story, the reader becomes aware of these different views in a particular order, and each one against a different context. The events in the main plot line – the brother rescues the princess – come as the first version presented, and the reader will accept them at face value. The events in the water carrier's false claim come later as an alteration of the truth that the reader has already accepted, so it is very likely that the reader will interpret this story as false in the face of that contradiction. The events told by the princess to her father the tsar arrive as a confirmation of the truth of the original version of the story, and as a denunciation of the falsity of the water carrier.

Second, because in the process of reading, the reader is aware that different characters, restricted in their perception, only have access to some of these versions and not others. The tsar, for instance, first hears the version of the water carrier, and very likely accepts it as the truth. The princess, on hearing the version of the water carrier, is already aware of the truth – having been involved in the actual

1	sets_out brother2
	start_story kidnapping
2	at_location palace kidnap dragon princess2 to_location dragons_lair
3	torment_at_night dragon princess2
4	tells_story someone brother2 kidnapping
5	decides_to_react brother2
6	sets_out brother2 to_location dragons_lair
7	fight brother2 dragon
8	wounded brother2
9	defeats brother2 dragon
10	releases brother2 princess2
11	sets_out princess2 to_location way_to_palace
12	at_location palace order tsar water_carrier bring_bones
13	sets_out water_carrier to_location way_to_palace
14	meets water_carrier princess2
15	returns [water_carrier+princess2] to_location palace
	start_story false_claim
16	at_location dragons_lair fight water_carrier dragon
17	defeats water_carrier dragon
18	releases water_carrier princess2
19	at_location palace tells_story water_carrier tsar false_claim
20	at_location dragons_lair sets_out brother2 to_location palace
21	arrives brother2 palace
22	recognises princess2 brother2
	start_story rescue
23	at_location dragons_lair fight brother2 dragon
24	wounded brother2
25	defeats brother2 dragon
	releases brother2 princess2
26	tells_story princess2 tsar rescue
27	exposed water_carrier
28	punished water_carrier
29	marry brother2 princess2

Table 6
Full discourse for the false hero sequence as described by Propp in his analysis of tale 155.

events – so she is unlikely to be deceived and very likely to recognise the lack of sincerity of the water carrier. Any reader will immediately recognise these facts. The proposed model intends to provide some factual data, extracted automatically from the discourse that may support this type of inference.

The proposed algorithm for constructing the set of views for different characters in a story traverses the discourse much as the previous algorithms described, but it builds separate representations for each of the characters. At each point in the discourse, it relies on the inferred model of the storyworld to predict which characters are capable of perceiving the events being reported, and updates the corresponding views only for those characters that are estimated to be present. During the processing of an embedded story, the events in the sub-story are also used to update the views of the characters known to be present at the telling of the embedded story. The outcome of this algorithm is a set of views,

	Preluding stories:
0	story **kidnapping**
1	kidnap dragon princess2
2	torment_at_night dragon princess2

	Main plot line:
1	sets_out brother2
2	tells_story someone brother2 **kidnapping**
3	decides_to_react brother2
4	sets _out brother2

	Anaphoric story by princess2		Events in main plot line		False claims by water_carrier
0	start_story **rescue**			0	start_story **false _claim**
1	fight brother2 dragon	5	fight brother2 dragon	1	fight water_carrier dragon
2	wounded brother2	6	wounded brother2		
3	defeats brother2 dragon	7	defeats brother2 dragon	2	defeats water_carrier dragon
4	releases brother2 princess2	8	releases brother2 princess2	3	releases water_carrier princess2
		9	sets_out princess2		
		10	orders tsar water_carrier bring_bones		
		11	sets_out water_carrier		
		12	meets water_carrier princess2		
		13	returns water_carrier		
			returns princess2		
		14	tells_story water_carrier tsar **false_claim**		
			START false claim plot line:		
			END false claim line		
		15	sets_out brother2		
		16	arrives brother2 palace		
		17	recognises princess2 brother2		
		18	tells_story princess2 tsar **rescue**		
			START anaphoric plot line:		
			END anaphoric plot line		
		20	punished water_carrier		
		21	marry brother2 princess2		

Table 7
Schematic representation of the plot lines of the various embedded stories told in tale 155. The main plot line is shown down the middle, the altered version told by the water carrier to justify his claim as false hero is shown on the right, and the anaphoric story told by the princess to refute it is shown on the left – both alligned with the events in the main plot for ease of comparison.

one for each character, such that each one includes only a part of the schematic that an omniscient narrator would construct for the complete story.

The current version of the algorithm is elementary in the sense that it cannot yet operate across different narrative levels: it will process the discourse for a narrative level and identify correctly the views of the characters expressed in that level. In its present form, it outputs a set of separate predictions: one for the main narrative and one for each of the substories involved in it.

The differences in perception between the different characters are significant. The set of views in itself is too verbose to report verbatim in a paper of this length. An approximate idea of the different degrees of coverage for each character can be obtained from Figure 1, which presents the part of the schematic for the story that each character has perceived as a coloured area over the schematic for an ominiscient narrator. For each character, the account of the views shown for them in this diagram has been compiled by hand over the predictions that the algorithm makes for the main narrative and each

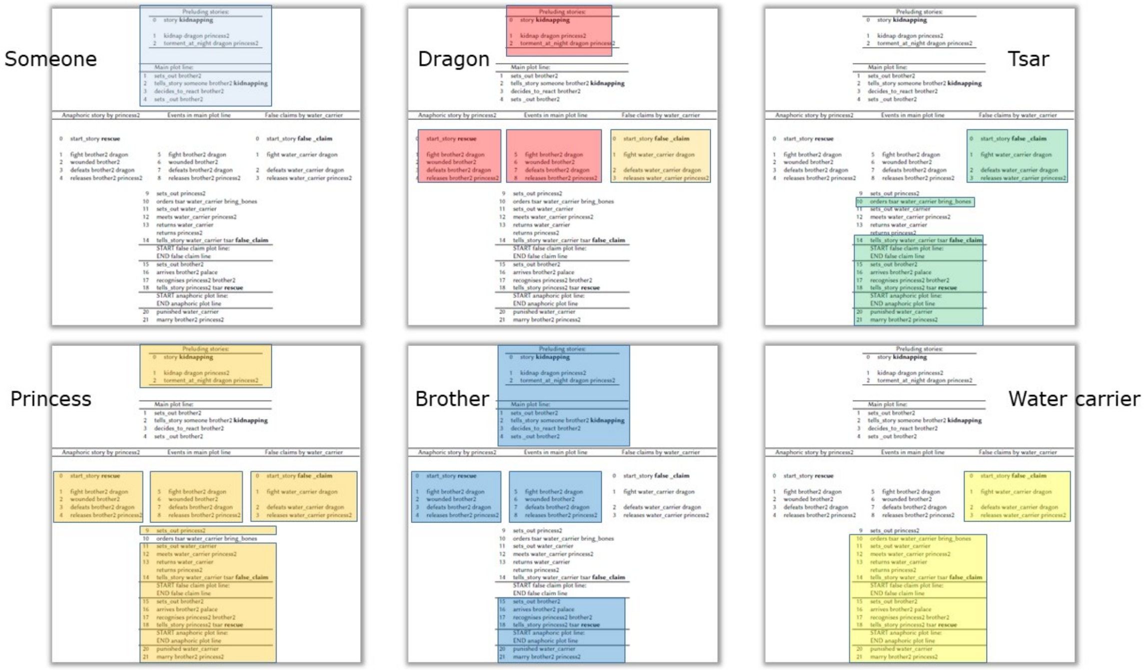

Figure 1: Colour coded coverage of the schematic for all the views in the story for the different characters.

of the substories involved in it.

The diagram shows that characters that remain at the same location throughout the story – the person who tells the hero about the kidnapping of the princess, the dragon and the tsar – have a much narrower view of the set of events in the story than the characters who travel across several locations – the princess, the hero and the water carrier.

The need to compile the views for different substories by hand raises some interesting questions regarding how the predictions of the view on the storyworld for characters present in an embedded story should be treated. For instance, a direct analysis of the story of the water carrier, where he falsely claims that he defeated the dragon and released the princess, predicts that the dragon, being mentioned as a participant, was present and is aware of the corresponding events. But we know the story to be false. So it would probably be fair to say that the dragon is unaware of these facts. In the diagram in Figure 1, the prediction arising from the insincere story told by the water carrier is shown in a lighter colour to indicate its uncertain status.

Similarly, when the princess hears this story, she is already aware of what actually transpired at her rescue, so she knows that the events described in the story did not transpire as told. However, in this case, she is aware that this version of the events has been put forward, so it is important that the model of her perception includes it. The interpretation machinery we are using allows us to tag this sub-story as false, which captures some of the relevant information.

4. Discussion

Discussion of the work presented is done in terms of shortcomings of the present version of the algorithm and in terms of relation to previous work.

4.1. Shortcomings of the Present Version

The algorithm as described in this paper has been implemented as a Java program that takes as input text files with the encoded version of discourse and produces text files that represent the structures shown in the tables in the paper. For the examples given, some effort has gone into transcribing these structures as readable LaTeX tables but their topology and information content has been respected.

The solution presented here operates at the level of a conceptual representation of discourse instead of accepting inputs in the form of text. This is a significant shortcoming in as much as it precludes the direct application of the procedure to real-world narratives, requiring an effort of hand-coding the corresponding texts in the notation for discourse representation used as input. However, it must be noted that the conceptual distinctions that the procedure is designed to identify – narrator, audience, embedded story – and the structures it is designed to build – narrative level, unrelated story, preceding story, anaphoric story, conflicting story, modal story – correspond to a level of abstraction that exists at a certain distance from the linguistic concepts required to interpret text in a particular language into the conceptual representation proposed as input. In this sense, the present proposal follows Ricoeur's dictum that research on the pragmatic use of language needs to be phrased at a level of abstraction beyond the grammatical devices of specific languages [10]. One may also note that Ryan's model of the modal nature of narrative universes [11] is also described over inputs consisting of propositions that are assigned a truth value.

The fact that the current version of the algorithm cannot yet operate across different narrative levels constitutes an important restriction. Where a story operates over more than one narrative level, the algorithm can be run at each level but the results across the various levels still need to be integrated by hand. However, the prototype as it stands has at least served to establish the relative importance of identifying relevant information such as who is narrating what to whom, when the discourse is switching to different narrative levels, and how different representations may be required for the different narrative levels of a story. The prototype has also demonstrated the feasibility of doing this automatically from a very limited conceptual transcription of the content of a story.

As an important insight arising from these observations, we identify the need for further research into the automated identification of boundaries between narrative levels in text, possibly phrased in terms of segmentation of the text into narrative levels, but with additional efforts to identify both narrator and audience for each particular narrative level. Efforts along these lines do exist, both in terms of attempts to construct knowledge resources that might support machine learning efforts along these lines [18] and actual attempts at carrying out the task automatically [19]. However, it is important to note that, regardless of whether the relevant information can be obtained automatically from text, the challenge of knowing what to do with the information for the adequate processing of narrative is still relevant. That challenge is what the present paper aims to address.

Once validated operational solutions are available for the task of automatically identifying narrative level data from text, they might be combined with solutions for transcribing textual content – sentences – as conceptual paraphrases of the kind employed as input in this paper – story updates. We assume that the type of solutions based on Large Language Model being developed these days [20] would not find excessive difficulty in the task of transcribing sentences as set of simple predicates with constants as arguments. If this combination were achieved, it would make it possible to apply the procedure proposed here to real-world narratives.

Regardless of the feasibility of this attempt to resolve the applicability of the proposed procedure to raw narrative text, the possibility of constructing automatically the data structures described for specific narrative levels still constitutes a valuable contribution. The insights obtained from the analysis of the views compiled by hand suggest that a set of simple heuristics based on the information that is available for each story may yield a reasonable baseline algorithm to address computationally the task of constructing fuller models that capture the relations across narrative levels. This has yet to be addressed as further work.

The role of embedded stories as means for providing characters who have not travelled with information about events that happened elsewhere is also captured by the algorithm. Thus, the hero learns about the kidnapping and the tsar learns about the rescue.

The work presented here raises some additional question that need to be considered. There are certain inferences that people will make automatically regarding the extent of the knowledge about the storyworld held by different characters that the algorithm in its present form still does not contemplate. For instance, the algorithm predicts that the tsar and the water carrier have no knowledge of the kidnapping, because they are not mentioned as being present when it happens and they are not present at the time that it is reported to the hero. However, common sense suggests that they are both aware of it. Similarly, the algorithm predicts that the hero has no knowledge of the version of the rescue presented by the water carrier, because he was not present at the moment when it was told. However, from the way the action develops it appears sensible to assume that, on arrival at the palace he is made aware of the water carrier's claim, even though that is never stated explicitly in the story. Consideration of how the algorithm might address this type of situation will be included in further work.

4.2. Relation to Previous Work

With respect to previous work reviewed in Section 2, the algorithm proposed captures the two main stages identified by psychologists as part of the process of interpretation of narrative: construction of a model of the physical world [2, 4] and construction of a model of how the different characters perceived the storyworld [1, 5]. It also provides a baseline model of the interpretation process for a narrative that includes features that may be engaged in the process of identifying material beyond a single reading of the story, which Barthes [7] considers an essential responsibility of the reader of narrative.

The machinery presented for identifying which parts of the story are perceived by each character would provide a valuable tool in any computational account of point of view as defined by Friedman [8]. The concept that Friedman describes goes beyond the idea of each character having a different view of the storyworld, as it involves not only each character perceiving different subsets of events in the storyworld but also giving each character a different approach to the interpretation of those events. Friedman's concept of point of view also includes the possibility of the author relying on her way of presenting the events to signal this different approach by each character to the way events are interpreted. These additional features are beyond the current capabilities of the model presented, but any model that intends to address them would need to rely on computation tools to partition the set of events in the story into the subsets that each character is aware of.

The technique of Character Perspective Charting proposed by Shanahan and Shanahan [9] supports the importance of interpreting stories from points of view other than those evident from a parse of the story as a whole, or even those of the main protagonist of the story. The algorithm proposed here would be a valuable tool to help identify the set of candidates to which Character Perspective Charting can be applied, and even to provide some baseline metrics that could help to select among them. From the example shown in Figure 1, it is clear that characters with larger coverage of the events in the story would be better candidates for treatment in terms of Character Perspective Charting. Additionally, one might consider that, in the case of stories having conflicting views on the storyworld, a selection that includes characters that subscribe to contrasting views might be better than one including characters that share the same views. The outputs of the proposed tool may automatically provide data to inform these decisions.

In terms of Ryan's model of the modal nature of narrative universes [11], the representation that we are proposing falls short of capturing the level of semantic detail that she contemplates. However, the representation does contemplate: a set of different possible worlds that must be updated during the interpretation of narrative discourse, the need to identify both the narrator and the audience at each point so the events being narrated can be used to update the appropriate possible worlds, and the consideration of modality to classify updates that affect possible worlds different from the actual world being built for the story.

The interpretation procedure described here matches Ryan's outline for her suggested algorithm in the sense that it traverses the input discourse to construct a data structure that describes at each point the partial outcome of "the act of reading/comprehending". Where Ryan considers sentences the implementation described here considers updates to the discourse. Where Ryan's model considers possible worlds, the implementation here considers separate branches in the interpretation graph. Where Ryan presents a conceptual description of her proposed algorithm, the description here corresponds to a working implementation. The implementation in question is far from the level of conceptual detail captured in Ryan's model, but it may be considered as a baseline approximation, intended mainly to establish the relative importance of addressing the problem in models of narrative interpretation.

5. Conclusions

The present paper describes a computational tool for extracting for a given story the set of restricted views on the events of the story that different characters may have. It operates on instances of narrative discourse that provide an outline of the story in terms of simple predicates grouped into conjunctive updates to the state of the story. The tool for extracting the set of views for different characters operates in two stages. First it creates a model of the physical world that compiles the set of locations mentioned in the world and infers from the discourse which characters are present at which location at particular moments in time. Then it relies on this model of the physical world to predict for each character which events in the story they are aware of.

The outcome of the tool is a set of views on the storyworld for the different characters in a story. The current version is limited to operating on single narrative levels. Over the resulting outcome, valuable information can be extracted on the coverage over the events in the story that each character has, and on which of the conflicting views on the story – if any are present – each character is aware of.

In general terms, the proposed tool models some of the operations of narrative interpretation that psychologists have considered relevant. It also provides a baseline solution to model the tasks that modern narratology considers fundamental in a reading process. The proposed model includes means to identify, represent and contrast different views on a story from the differences in physical perception of the world by different characters.

As future work, we will consider the extension of the algorithm for identifying character views to be able to operate correctly across more than one narrative level. The tool as it stand provides the raw material to consider – the set of predictions on the views of characters for each sub-story – and information that would be relevant to the decisions needed to integrate them - such as relative truth values on the sub-stories or awareness of the different character of each of the sub-stories.

We hope to test the algorithm over a set of different stories to ensure robustness in the face of a range of linguistic and pragmatic situations. As part of this effort we are specially interested in testing the algorithm on stories that involve several narrative threads taking place in parallel at different locations. We reckon that the outcomes on this type of input may provide more useful insights than on simple stories.

Acknowledgments

This paper has been partially funded by the projects CANTOR: Automated Composition of Personal Narratives as an aid for Occupational Therapy based on Reminescence, Grant. No. PID2019-108927RB-I00 (Spanish Ministry of Science and Innovation), the ADARVE (Análisis de Datos de Realidad Virtual para Emergencias Radiológicas) Project funded by the Spanish Consejo de Seguridad Nuclear (CSN) and project DARK NITE: Dialogue Agents Relying on Knowledge-Neural hybrids for Interactive Training Environments, Grant No. PID2023-146308OB-I00 (Spanish Ministry of Science and Innovation).

References

[1] N. Tal-Or, J. Cohen, Understanding audience involvement: Conceptualizing and manipulating identification and transportation, Poetics 38 (2010) 402–418. doi:10.1016/j.poetic.2010.05.004.

[2] J. Brunner, Actual Minds, Possible Worlds, Harvard University Press, 1986. URL: https://books.google.es/books?id=--_ItAEACAAJ.

[3] P. Gervás, A discourse interpretation engine sensitive to truth revisions in a story, in: Tenth Annual Conference on Advances in Cognitive Systems, Arlington, Virginia, 2022.

[4] R. J. Gerrig, Experiencing Narrative Worlds: On the Psychological Activities of Reading, Yale University Press, 1993. URL: http://www.jstor.org/stable/j.ctt1dr37cb.
[5] B. Teasdale, L. Maguire, F. Budelmann, R. I. M. Dunbar, How audiences engage with drama: Identification, attribution and moral approval, Frontiers in Psychology 12 (2021). doi:10.3389/fpsyg.2021.762011.
[6] J.-P. Vernant, P. Vidal-Naquet, Mythe et tragédie en Grèce ancienne, volume 2, La Découverte, 1972.
[7] R. Barthes, The death of the author, 1968, na, 1992.
[8] N. Friedman, Point of view in fiction: the development of a critical concept, PMlA 70 (1955) 1160–1184.
[9] T. Shanahan, S. Shanahan, Character perspective charting: Helping children to develop a more complete conception of story, The Reading Teacher 50 (1997) 668–677.
[10] P. Ricoeur, Interpretation theory: Discourse and the surplus of meaning, TCU press, 1976.
[11] M.-L. Ryan, The modal structure of narrative universes, Poetics Today 6 (1985) 717–755.
[12] D. Herman, M. Jahn, M. Ryan, Routledge Encyclopedia of Narrative Theory, Taylor & Francis, 2010.
[13] J.-K. Adams, Order and narrative, Recent Trends in Narratological Research (1999) 111–27.
[14] P. Gervás, A model of interpretation of embedded stories, in: Text2Story: 4th International Workshop on Narrative Extraction from Texts, CEUR Workshop Proceedings, CEUR Workshop Proceedings, Lucca, Tuscany, 2021.
[15] V. I. Propp, Morphology of the Folktale, volume 9, University of Texas Press, 1968.
[16] D. Wilson, D. Sperber, On grice's theory of conversation, in: Conversation and discourse, Routledge, 2022, pp. 155–178.
[17] G. Genette, J. Lewin, J. Culler, Narrative Discourse: An Essay in Method, Cornell paperbacks, Cornell University Press, 1983.
[18] E. Gius, J. Jacke, The hermeneutic profit of annotation: On preventing and fostering disagreement in literary analysis, International Journal of Humanities and Arts Computing 11 (2017) 233–254.
[19] N. Reiter, J. Sieker, S. Guhr, E. Gius, S. Zarrieß, Exploring text recombination for automatic narrative level detection, in: Proceedings of the Thirteenth Language Resources and Evaluation Conference, 2022, pp. 3346–3353.
[20] Z. Li, Y. Zeng, Y. Zuo, W. Ren, W. Liu, M. Su, Y. Guo, Y. Liu, L. Lixiang, Z. Hu, et al., Knowcoder: Coding structured knowledge into llms for universal information extraction, in: Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), 2024, pp. 8758–8779.