Interpreting Narrations of Events Witnessed: Relying on Location Data to Help Place Embedded Stories

Pablo Gervás[1]

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

Abstract

When interpreting a narrative in which someone tells a story, first-person accounts of events that happened to the narrator earlier are relatively easy to process. But cases in which the narrator is telling something that she witnessed – but was not personally involved in – require more elaborate pragmatic inferences based on the assumption that the narrator was present as a witness of the events being told. Appropriate handling of these cases requires means of correct interpretation of where events experienced and events witnessed – by the narrator – are happening in terms of locations in the storyworld. The present paper considers a simple model of how information on changes in location across events may be inferred from narrative discourse, and how this information can be exploited to ensure appropriate temporal placement – with respect to the events of the main or frame story – of the events for the embedded stories that were witnessed by a narrator at some earlier point in the discourse.

Keywords

interpretation of narrative discourse, embedded stories, changes in location, inferences based on discourse, narrator as witness

1. Introduction

When we consider a story as a stand alone product we tend to disregard the context in which it is produced: the story is usually an act of communication between a narrator and an audience [1], but both of these are implicit in the existence of the story. In contrast, narratives in which one characters tells a story to other characters provide explicit information not only on the story, but also on the context in which it occurs. In these cases, the main narrative – in which a character is telling a story – is known as the *frame story*, and the story being told by the character is known as the *embedded story*. In that sense, they constitute a very valuable source of information on how narrative operates as a communication device. A reader faced with such narratives needs to process the embedded story in the appropriate context, not only to understand the embedded story itself, but also to extract from it the relevant information that may influence their understanding of the frame story. This involves not only identifying them as stories told by the characters in a frame story, but also establishing how the events in the embedded story relate to the events in the frame story. With respect to the frame story, events in the embedded story may happen in a different world, be a retelling of parts of the frame story that have already happened, or be part of an alternative telling of parts of the frame story [2]. To inform the task of correct placement, existing solutions for the interpretation of embedded stories rely on comparisons between the characters and actions involved in sub-sequences of the frame story and the embedded story. This type of comparison provides valid solutions for cases in which the narrator is involved in the embedded story that she is telling, because her very presence in the events in the embedded story provides the required information for a correct placement. However, there are cases in which the narrator is not directly involved in the events that she is telling about, and yet human readers can correctly interpret how the embedded story fits into the frame story based on pragmatic assumptions

In: R. Campos, A. Jorge, A. Jatowt, S. Bhatia, M. Litvak (eds.): Proceedings of the Text2Story'25 Workshop, Lucca (Italy), 10-April-2025.

*Corresponding author.
pgervas@ucm.es (P. Gervás)
http://nil.fdi.ucm.es// (P. Gervás)
0000-0003-4906-9837 (P. Gervás)

about similarity between the locations in which the narrated events happen and locations that we know the narrator has recently visited in the course of the frame story. The present paper describes a simple model that combines inferences about embedded stories and inferences about changes in location to deliver a baseline model of how these more complex cases might be handled.

2. Previous Work

Three topics are reviewed to inform the work reported: basic challenges of reading narrative that involves embedded stories, how to place embedded stories with respect to the frame story and how to infer changes in location from narrative discourse.

2.1. Reading Narrative with Embedded Stories

Psychologists studying what takes place in the mind of a reader processing fiction have identified that a fundamental part of the process is the construction of a physical model of the fictional world in which the story takes place [3, 4]. *Embedded stories* appear in narratives where a certain character tells a story to other characters [5]. The story in which the telling of the story takes place is known as the *frame story* and the story that is being told is known as the *embedded story*. This type of situation where one story is told within another is said to involve two different *narrative levels*: an outer one for the frame story and an inner one for the embedded story. In the case of stories created by an author for a wider audience, a basic narrative level is established that involves the author as narrator and the audience as intended recipient.

Todorov [6] and Tenev [7] both identify modality as a fundamental ingredient in the construction of narrative, with conflict in stories often arising from contrast between actual and desired situations or between opposing obligations/desires. Ryan [8] elaborates on this idea, describing narrative as a sequence of moves that transition between states of the storyworld that include models of characters views on wishes, obligations or beliefs. In both cases, the ability to model this kind of evolution on the modal views of characters becomes a fundamental aspect for narrative interpretation. As the views of characters on wishes, obligations or beliefs are often presented as embedded stories [2], the correct interpretation of embedded stories becomes relevant for correct identification of plot.

2.2. Relative Placement of Embedded Stories

Gervás [9] proposed an algorithm for the interpretation of embedded stories that relied on on a stack-based mechanism for handling the changing contexts of interpretation. Rather than work on text, this algorithm operated on a simplified format for conceptual description of narrative discourse designed to capture the information relevant for testing hypotheses on the operations and data structures required for adequate processing of embedded stories. The discourse for a story is therefore represented as a list of updates to the story, where an update is a conjunction of statements – each a predicate with arguments – that jointly describe the event. The sequence of updates to the story describes the sequence of events or facts for the story. Specific statements are used to indicate the start of an embedded story and the fact of its telling with the frame story, with the sequence of statements for the embedded story appearing bracketed between them in the sequence for the discourse of the frame story. The algorithm allowed the correct separation between narrative levels and the identification of the spans of events in the discourse corresponding to embedded stories.

An extension of this algorithm [2] introduced mechanisms for estimating the relative placement – within a partial reconstruction of the fabula for the story – of the spans of discourse corresponding to embedded stories with respect to the discourse for the frame story. This extension relied on comparing – for the pairs of spans under consideration from frame story and embedded story – the sets of characters involved and the sequence of events described. As the work presented in this paper further extends these solutions, the basic ideas of each one of them – as required to understand the proposal made here – are briefly described below for ease of reference.

Table 1
Basic algorithm for interpretting embedded stories.

- 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

The algorithm for the interpretation of embedded stories described in can be summarised in Table 1. This results in a *story interpretation* for the frame story – with telling predicates to indicate where embedded stories are told – and a table of (*story interpretations* of) embedded stories indexed by name. 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 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*). 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.

The algorithm proposed for this in [9] operates by constructing a *branching partially ordered graph* that compiles the events and sub-stories in the discourse according to estimates of those relations determined by basic heuristics. The procedure proposed for building this branching partially ordered graph is shown in Table 2. The procedure for establishing whether an embedded story partially matches a given span of discourse relies on alligning updates in both spans whenever the corresponding conjunctions of statements have matching statements. Statements are considered to match if they share the same predicate name and at least the value for one argument (in the same argument position of the predicate). Whenever more than one match is found, the final match is selected based on length of matching span and average value of percentage of shared arguments across all the matched statements.

Table 2
Algorithm for placing fabulae for embedded stories with respect to fabula from frame story.

- 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)

2.3. Inferring Shifts in Location from Narrative Discourse

A related effort to model features involved in the interpretation of narrative discourse considers how a physical model of the world may be built from a discourse [10]. This approach relies on a two-step process: first the explicit information available in the discourse on storyworld locations and location changes is compiled into an intermediate data structure, then a set of heuristics is applied to expand this data structure into a fuller representation of the locations in the storyworld. The heuristics in question rely on basic assumptions about continued existence of locations in the storyworld if they are mentioned at any point in the discourse and continued presence of characters at a given location unless they have been explicitly described as moving away.

Table 3 shows an example of the process of interpretation of the information on locations and character presence as extracted by the algorithm. An example of story is given at the top, showing both a text summary and the discourse representation employed here. This story corresponds to an excerpt from tale 155, as analysed by Propp in the appendix of this book [11]. The two steps of world model construction are shown below. 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: brother to dragon's lair, princess away when she is released.

3. Informing Placement of Embedded Stories with Location Shift Information

The stories we are interested in involve situations in which one of the characters tells stories about events he has witnessed, but in which he was not involved. Although the narrator may mention explicitly in his story what he was doing at the time, this does not necessarily happen. When there are no such mentions, embedded stories of this type cannot be associated to the frame story by relying on the simple heuristic of identifying shared characters across frame and embedded story. Table 4 shows examples of such a story, presented as text in column (a) and as the full transcription of the conceptual representation of the discourse used as input in column (b).[1]

When this story is processed with the heuristics applied in [9], the two embedded stories are identified as preluding stories (because they take place at locations mentioned in the frame story, they are considered related) but the relative chronology of each one of them with respect to the frame story and the relative chronology between them cannot be identified. An example of this imperfect

[1]Although I would have preferred to use examples of real-world narrative, the example used here has been engineered to include the maximum set of problematic features in the minimum space that would allow for the resulting output structures to fit in the available space. As mentioned later in Sections 4 and 5, further work should consider the application of the procedures described here to real-world narratives.

Table 3
Examples of a basic story involving changes of location, given as text (a) and transcribed into the discourse representations employed here (b); followed by the world model (c) built in two steps for an input narrative discourse. In this world model, the first column shows the discourse representation used as input – with additional predicates describing location-related information – the next column (Step 1) shows explicit location information as compiled in the first step, and the final column (Step 2) shows enriched model of all locations in the storyworld over time, inferred via heuristics.

(...) The second brother sets out to rescue a princess kidnapped by a dragon. He fights the dragon and receives a wound in the process. He defeats the dragon. He liberates the princess. (...)

(a) text for the story

1	set_out brother destination dragonlair	4	defeat brother dragon
		5	release brother princess
2	fight brother dragon	6	set_out princess destination palace
3	wounded brother		

(b) input discourse representation

	Step 1: **WorldSnapshot**		Step 2: **Inferred Representation**		
Story updates	**InitialLoc**	**Destination**	**loc-1**	**dragonlair**	**palace**
`<default context>`	**loc-1** [brother]		[brother]	[dragon, princess]	[]
set_out brother **destination dragonlair**	**loc-1** []	**dragonlair** [brother]	[]	[dragon, princess, brother]	[]
fight brother dragon	**dragonlair** [brother,dragon]		[]	[dragon, princess, brother]	[]
wounded brother	**dragonlair** [brother]		[]	[dragon, princess, brother]	[]
defeat brother dragon	**dragonlair** [brother,dragon]		[]	[dragon, princess, brother]	[]
release brother princess	**dragonlair** [brother,princess]		[]	[dragon, princess, brother]	[]
set_out princess **destination palace**	**dragonlair** []	**palace** [princess]	[]	[dragon, brother]	[princess]

(c) world model construction steps

reconstruction of the fabula is shown in column (a) of Table 5. This presents a schematic representation of the fabula reconstructed for John's story by the original algorithm.

The solution proposed in this paper is to consider the models of the storyworld constructed for each of the stories (one for the frame story and one for each of the embedded stories) to inform the process of determining the relative placement in time of the embedded stories with respect to the frame story.

To achieve this, we rely on the assumption that when the protagonist of a story such as this one tells other characters about events that have happened at locations that he is visited during the story, he is most probably describing events that he was witnessed while he was present at those locations.

With respect to the algorithm described in Section 2.2, the following modifications are required to achieve this:

- if an embedded story is found to be related to the frame story (not an unrelated story), but no match is found with the frame story (not a conflicting story)
- the world model for the embedded story is checked for the set of locations at which it takes place
- the narrator of the embedded story is identified
- the world model for the frame story is checked for time points at which the narrator of the embedded story was present at the locations in which the embedded story takes place
- the embedded story is assumed to have taken place between the start and the end of the period during which the narrator was at those locations
- the embedded story is placed in the graph for the fabula at the estimated time point

An example for the graph obtained for the fabula by applying this procedure is shown in in column (b) of Table 5. This shows a schematic representation of the fabula reconstructed for John's story by the

Table 4
Examples of a story of the type under consideration, given as text (a) and transcribed into the discourse representations employed here.

John woke up at home and had breakfast. Then he went to school. He had a maths class in the maths classroom. He had lunch. He had an English class in the English classroom. He went to the soccer field to have soccer practice. Then he returned home. When he got home, he told his mother how Peter had released his pet rat in Maths class, the pet rat scared the teacher and Peter got punished. Then he told his mother how Mike had an accident during soccer practice, and how the coach helped Mike, who had to go to hospital.

(a) text for the story

1	at_location home wakes_up john	10	start_story peter_incident
2	has john breakfast	11	releases peter pet_rat
3	sets_out john to_location school	12	fears teacher
4	at_location maths_classroom has john maths_class	13	punished peter
5	has john lunch	14	tells_a_story john mum peter_incident
6	at_location english_classroom has john english_class	15	start_story mike_incident
7	sets_out john to_location soccer_field	16	at_location soccer_field had_accident mike
8	has john soccer_practice	17	decide_to_help coach mike
9	sets_out john to_location home	18	sets_out mike to_location hospital
		19	tells_a_story john mum mike_incident

(b) input discourse representation

extended version of the algorithm, enhanced with heuristics for identifying parallelisms between embedded stories and spans of the frame story based on location information inferred from the storyworld models for frame and embedded stories.

4. Discussion

The contrasting representations – (a) and (b) – presented in Table 5 for the input story shown in Table 4 illustrate how the proposed extension to the narrative interpretation algorithm allows for more accurate representation of the interpretation that intuitively occurs to human readers. This is not surprising as the proposed extension enriches the previous version of the algorithm with a process of explicit construction of a model of the storyworld, known to be a fundamental element in human interpretation of narrative [3, 4].

An important insight arising from the work presented is that such a model of the physical aspects of the storyworld has an importance in the interpretation process beyond the simple acquisition of information about the world: without such a physical model of the storyworld inferences on relative placement of embedded stories may be incorrect or underdetermined.

The algorithm in its present form 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. Some effort has gone into transcribing these structures as readable LaTeX tables but their topology and information content has been respected.

The algorithm as described should be considered as a baseline approximation, intended mainly to establish the relative importance of addressing the problem in models of narrative interpretation. The current version is susceptible to incorrect assumptions as to which moments of a character's past correspond to the point in which a given embedded story is witnessed. More refined heuristics to inform such decisions may be considered as further work.

An additional issue arises from the fact that the process of constructing the physical model of the storyworld happens in two stages, one of recording explicitly provided information on location and one of filling in missing information about the world based on heuristic-driven inference. Because the

Table 5
Representation interpreted by the original algorithm (a) and the revised version (b).

Preceding stories:	
Peter's incident	Mike's incident
start_story peter_incident	start_story mike_incident
at_location maths_class	at_location soccer_field
releases peter pet_rat	had_accident mike
fears teacher	decide_to_help coach mike
punished peter	sets_out mike
	to_location hospital

Main plot line:
at_location home wakes_up john
has john breakfast
sets_out john to_location school
at_location maths_class has john maths_class
has john lunch
at_location english_classroom has john english_class
sets_out john to_location soccer_field
has john soccer_practice
sets_out john to_location home
tell_story john mum peter_incident
tell_story john mum mike_incident

(a) interpretation by original algorithm

Main plot line:
at_location home wakes_up john
has john breakfast
sets_out john to_location school

Main story + side story	
John's maths class	Peter's incident
at_location maths_class has john maths_class	start_story peter_incident
	at_location maths_class releases peter pet_rat
	fears teacher
	punished peter

Main plot line (continued I):
has john lunch
at_location english_classroom has john english_class
sets_out john to_location soccer_field

Main story + side story	
John's soccer practice	Mike's incident
has john soccer_practice	start_story mike_incident
	at_location soccer_field had_accident mike
	decide_to_help coach mike
	sets_out mike to_location hospital

Main plot line (continued II):
sets_out john to_location home
tell_story john mum peter_incident
tell_story john mum mike_incident

(b) interpretation by revised algorithm

second step must always be defeasible – in the sense that if information provided later in the discourse contradicts any of the inferences made, those inferences need to be revised – this implies that if the placement of any embedded story has been based on such defeasible inference, it will itself be defeasible. The procedure described here may need to be reconsidered in these terms in further work.

The introduction of this type of defeasibility into the computational process opens up possibilities for the author wishing to exploit it for particular effects. An author may know the type of inferences that a reader may be making in processing each possible instantiation of the discourse, in terms of default assumptions as to which character is where, and therefore also about which character is aware of which events in the storyworld. By exploiting this type of information, such an author may craft particular discourses that, without being openly misleading, may result in the reader making assumptions that will obfuscate the parts of the plot that the author wants kept in the dark.

The work described in this paper is limited to a conceptual description of discourse restricted to a small subset of the features that may be expressed in natural language. In particular, it focuses on representation of events and location information of where they take place. More elaborate solutions to the problem of relative temporal placement of the events in embedded stories with respect to the frame story will be required once additional contextual cues – such as temporal expressions or discourse markers – are considered.

In more pragmatic terms, future work should consider identification of appropriate machine learning solutions to replace the hand-coded heuristics currently powering the implementation. However, in order to train such machine learning solutions a corpus of examples annotated with all the relevant data would be required. If the compilation of such a corpus were to be carried out entirely by hand, it would involve a very significant effort of knowledge engineering in itself. To simplify this process, we proposed a three stage procedure for bootstrapping such a corpus. An initial stage would rely on state of the art technologies for natural language processing to construct a corpus of examples of real world narratives transcribed into the conceptual representation used in this paper. We assume that solutions based on Large Language Model would not find the task onerous. The existing implementation of the algorithm would then be used to construct examples of outputs that capture the complex structure of the narratives in terms of embedded stories and relative temporal placement of the corresponding fabulae. These outputs may then be revised by human experts to ensure correctness. This step would allow repair of any errors resulting from the use of limited heuristics. This would result in two parallel corpora of the initial texts paired with the corresponding representation for their underlying narrative structure. Based on such the corresponding parallel corpus, machine learning solutions for the task may be trained.

5. Conclusions

The present paper argues the importance of including a construction of the physical model of the storyworld in any process of narrative interpretation, and how such a model can help inform the correct temporal placement of embedded stories with respect to the fabula for the frame story in cases where the narrator of the embedded story is merely a witness of the events in the story. The algorithm proposed to address this task is a simple baseline intended to underline the importance of the task within the broader context of models of the interpretation of narrative. Additional research is needed in terms of both more refined heuristics for the task and further experiments with a wider range of examples of input stories.

The model as described has implications for automated storytelling, AI-based narrative understanding and discourse analysis. Automated story tellers may need to consider the use of embedded stories as part of the toolkit available to describe their storyworlds. Efforts at narrative understanding need to ensure that their solutions can handle embedded stories. Approaches to discourse analysis must include a description of discourses involving more than one narrative level.

In terms of further work, there is a clear need for additional empirical evaluation and computational experiments over a broad range of real-world narrative datasets. The functionalities outlined in this paper should be tested over different narrative styles or storytelling conventions. The pragmatic considerations presented at the end of Section 4 may provide a possible way forward in this sense. Once the approach is taken beyond simple explanatory examples as the ones presented in this paper, issues will arise with more complicated narratives, for instance, having situations in which multiple embedded stories conflict or overlap in terms of location and chronology. The extensions based on machine learning contemplated in Section 4 may fare much better than hand-crafted heuristics in dealing with such complex cases.

Acknowledgments

This paper has been partially supported by project 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), 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) and project ADARVE (Análisis de Datos de Realidad Virtual para Emergencias Radiológicas) funded by the Spanish Consejo de Seguridad Nuclear (CSN).

References

[1] P. Ricoeur, Interpretation theory: Discourse and the surplus of meaning, TCU press, 1976.
[2] 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.
[3] J. Brunner, Actual Minds, Possible Worlds, Harvard University Press, 1986. URL: https://books.google.es/books?id=--_ItAEACAAJ.
[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] D. Herman, M. Jahn, M. Ryan, Routledge Encyclopedia of Narrative Theory, Taylor & Francis, 2010.
[6] T. Todorov, Poétique de la prose, Média Diffusion, 2014.
[7] D. Tenev, Point of view and the modalities of narrative, Prace Filologiczne. Literaturoznawstwo [PFLIT] 1 (2018) 27–42.
[8] M.-L. Ryan, The modal structure of narrative universes, Poetics Today 6 (1985) 717–755.
[9] 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.
[10] P. Gervás, Basics of narrative interpretation: Physical model and character-specific views of the storyworld, in: 8th International Workshop on Computational Models of Narrative (CMN'25), Geneva, Switzerland, 2025.
[11] V. I. Propp, Morphology of the Folktale, volume 9, University of Texas Press, 1968.