Advances in Cognitive Systems X (20XX) 1-6 Submitted X/20XX; published X/20XX

A Discourse Interpretation Engine Sensitive to Truth Revisions in a Story

Pablo Gervás PGERVAS@UCM.ES
Facultad de Informática, Universidad Complutense de Madrid, Madrid, 28040 Spain

Abstract

The discourse for a narrative often conveys at different points of its span conflicting views on the events it is describing. This is because the discourse is intended to reflect not the truth of those events at a given point in time, but the evolution of the knowledge about those events in the experience of the narrator – or rather the evolution that the narrator wants the reader to experience. This implies that the discourse for a complete narrative will usually include several conflicting views in terms of what events are true. A process of interpretation of the story would need to model how these various views are constructed, stored, and validated or falsified as the reader progresses through the reading of discourse. The present paper presents a computational model of the cognitive process that might be required to accomplish this task.

1. Introduction

When using a linear discourse to describe a a particular state in a fictional world, each statement we make conveys details that we want to be true in the world we are creating. A reader processing our discourse creates a vision of the story world in their mind, adding detail from each statement. But some of the statements we make should not be processed in that way. When we decribe a plan, a wish, a promise…the events mentioned in them must not be added to the story world as being directly true. They describe situations that are desired by characters in the story – in different degrees – but which are not true in the story world at that point. These discourses usually occur as arguments of statements about what someone plans, wishes, or promises. We refer to them as *embedded discourses*. Yet some types of embedded discourse do refer to the story world under construction. When a character tells about what happened to them, the events they mention are assumed to have been true. If these events have already been told as part of the preceding discourse of the story, the telling may be interpreted as a reference to them. If they had not been told before, a reader processing such statement does add the events to their vision of the story world as believed to be true. This mechanism is often exploited maliciously by come characters, to mislead others by telling them false stories about what has happened. When this occurs, it triggers the construction of two differents interpretations of the story world: one where the original (real) story is true, and one where the told (false) version of the story is true. Authors often resort to this mechanism to create conflict and uncertainty in their plots.

The present paper proposes a computational model for the interpretation of embedded discourses, and explores the consequences and challenges of this behaviour for the overall task of interpretation of narrative. The model covers the three different uses described above: modal stories, anaphoric stories and false stories.

2. Previous Work

To provide the background for the work that is described in the paper, a number of topics need to be covered: existing models of how discourse is processed, some insights from psychology on how people process fiction and existing representations and computational models of dealing with narrative features.

2.1 Modelling the Interpretation of Discourse

Paul Ricoeur (1976) observed that much of the surge of very valuable work of modern linguistics – at the time – had focused on "the particular structure of the particular linguistic system", leaving the problem of discourse – the study of how language is used to achieve particular tasks – to "recede from the forefront of concern and become a residual problem". Since his time, the advent of computational linguistics has made very significant advances, but the trend to focus on syntax and semantics to the detriment of discourse as used for specific purposes has remained.

Indeed, many of the computational solutions developed over the intervening years for processing discourse focus on the interpretation of how the elements of a particular language combine to create logically represented meaning, with little effort devoted to the way in which the resulting meaning is applied to practical tasks. Such is the case for *Discourse Representation Theory* (Kamp & Reyle, 1993) or *Segmented Discourse Representation Theory* (Lascarides & Asher, 2007), which provide very good procedures for converting language statements into conceptual representations based on first order logic enriched with structures to account for phenomena such as co-reference. The type of co-reference they consider, however, is restricted to reference to entities, and reference to sequences of events described previously is not considered.

The present paper attempts to shed light on how discourse – considered as independent from a particular linguistic system – is employed in a nested fashion either to allow anaphoric reference to sequences of events described previously or as means of presenting alternative versions to to sequences of events described previously. Ricoeur's view on discourse pictures it as a sequence of sentences, where each sentence involves a predicate applied to some entities that need to be identified by the subject (and objects) of the sentence. He claims that (a sentence in a) discourse is the synthesis of two functions: the identification of what is talked about and the predication of something about it. He distinguishes between these functions and the set of "grammatical devices" which serve each of those functions – which would correspond to the syntax of a particular language. In this paper we will distance ourselves from particular natural languages by establishing a representation language that includes the minimally required features for identifying referents – constants – and for predicating upon them – predicates.

Ricoeur also points out that, discourse – being an act of communication – connects an event of speaking – or writing – and an act of hearing – or reading. Of particular interest for the present paper is his observation that, given that our statements may sometimes be ambiguous, discourse requires processes to "reduce the plurality of possible interpretations, the ambiguity of discourse" (1976, page 17). It is such processes, specific to ambiguity in the truth value assigned to some embedded stories, that we intend to model in this paper. In particular, we will be looking at: (1) how the interpretation of a discourse involving embedded stories – which allow telling at one point in the discourse about events in the past – may introduce ambiguity with respect to how the same events were told at an earlier point in the discourse, (2) how the resulting ambiguity in the truth value attributed to such events is handled during interpretation, and (3) how the representation for the accumulated ambiguities evolves over time – usually towards a resolution of the ambiguity as described by Ricoeur.

2.2 How People Process Fiction

From the point of view of psychology, there have been studies of the way in which people process fictional stories. For the purpose of clarifying the challenge addressed in this paper, it is important to keep in mind the definition of narrative proposed by Labov (1972): "a method of recapitulating past experience by matching a verbal sequence of clauses to the sequence of events which actually occurred". This definition brings together the two fundamental elements that need to be identified: the sequence of clauses that constitutes the discourse for the narrative and the set of events that are being conveyed. The task of interpreting a narrative discourse involves being able to reconstruct the set of events – and their relative temporal order – from the discourse sequence. Gerrig (1993) explores the cognitive operations involved in the task of reading fiction, with particular attention to the view in which this implies and exploration of the narrative worlds of books and movies. He mentions that, if the reader is to be able to make inferences on what she is reading, the process of interpreting a story must involve the construction of (something like) *situation models* (Van Dijk & Kintsch, 1983), which "integrate information from the text with broader real-world knowledge". These situation models represent in the reader's mind the world in which the story takes place. He proposes two metaphors to explore this challenge: one of being transported into these worlds, and one of the reader performing the stories. In the first one the reader reacts as if placed herself in that situation. In the second one the reader reacts as an actor asked to perform the actions affecting the characters. Under each of these views, the reader needs to reconstruct from the discourse of the narrative not just a single description of the world, but rather a sequence of descriptions that represent the states of the world as the story evolves. This is because, as a result of the plot, the situation in the story world keeps changing.

This is particularly noticeably in the analysis of the role of the reader in the reading task as carried out by Fish (1980). In this analysis, Fish makes special emphasis on the fact that the meaning constructed is not a single unit that describes the reader's view of the narrative, but rather the meaning is made up of the moment-by-moment account of the sequence of mental operations that the reader had to perform as part of the interpretation task. The readers' observation of their own performance during interpretation becomes a significant contribution to their experience of the narrative.

Brunner (1986) explores the mental operation involves in the construction of imaginary worlds within our minds. This particular task is a fundamental element in the performance of the reader:

as the narrative is interpreted, progressively expanding versions of the story world are being built by the reader. In fact, representations may need to be built from more than one such story world. Teasdale et al. (2021) carry out a number of experiments to explore people in an audience identify with characters, develop a moral judgment of their behaviour, or understand why they do what they do. In this process, it is clear that people construct models of the story world that are different and separate from the model they have of the real world. Furthermore, the fundamental process of mentalization – the ability to understand the mindstates of the storyteller and the character in the story (Tal-Or & Cohen, 2010) – requires the construction of separate models of what each character in the story, at each point of the story, considers to be true in the story world.

2.3 Embedded Stories and Narrative Order

Two topics from the theory of narratology need to be described: embedded stories – because they constitute instances of embedded discourse – and narrative order – because embedded discourses are often used to convey short stories about characters that are presented out of the chronological order of the main story.

Although "embeddings" have recently acquired fame in relation to neural representations of language, before that they were used in narratology to describe 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" (Herman et al., 2010). Such devices introduce the concept of a *narrative level* (Genette et al., 1983), which, for a story being told, describes the relative position at which a narrator and their audience stand with respect to the story – usually outside the story. Whenever characters within the story tell a story themselves, this constitutes an *embedded story* within the original story, now also referred to as the *frame story*. This nesting of embedded stories within overarching frame stories can occur repeately in a recursive manner.

Embedded stories and the relation between narrative levels have not traditionally been included in models for representing relations over a discourse, such as the original *Rhethorical Structure Theory* (Mann & Thompson, 1988) nor in later revisions (Taboada & Mann, 2006). More recent solutions specifically focused on narrative still address relations internal to a given story rather than the relations between a story and the context in which it is told (Nakasone & Ishizuka, 2006). Embedded stories are again not included among the set of features of discourse modeled by *Discourse Representation Theory* (Kamp & Reyle, 1993) nor *Segmented Discourse Representation Theory* (Lascarides & Asher, 2007).

It is only very recently that the computational processes required for handling embedded stories and narrative levels have been addressed in (Gervás, 2021b). In that paper, the discourse for a story is represented using a simple formal language based on predicate-argument structure that defines specific statement formats for declaring the start and end of embedded stories. A set of data structures is proposed to represent the discourse for stories that have more than one narrative level, and a stack-based interpretation procedure is defined that handles an input linear discourse with embedded stories by relying on a stack of interpretations to address the potentially recursive nature of embedded story telling.

An important aspect of narratology that is very relevant for the present paper concerns the concepts of *narrative order* – defined as a structured order to tell events which is independent of their dates – and *anachrony*, defined as the "the contrast between chronological and textual order" (Genette et al., 1983). It is very typical for the order of presentation of events in a narrative to not correspond to the order in which the event actually happened in the world, giving two different orders of sequencing for the events in a narrative: the chronological order in which they happened and the order in which they appear in the discourse for the narrative (Adams, 1999). Of particular interest to the present paper is the idea – expressed by Adams – that the events in embedded stories tend to break the expectation that events in a narrative appear in something close to the chronological order in which they happened. Adams explains that this is because embedded stories often act as supplemental explanations, providing – out of chronological order – information about events that have occurred earlier. Another important point covered by Adams is that there is a certain impossibility to determine – from the discourse of a story – the absolute or purely chronological order of events, but that being able to identify such an order is not actually important to understand – and benefit from – the narrative. We will take advantage of this concept in this paper to establish heuristics that identify a possible order between events mentioned in a discourse, without pretending to identify the real order in which they occurred in the world, but rather to propose a plausible order that is compatible both with the logical relations between them and with the way they are presented in the discourse.

3. Developing an Interpretation Engine for Narrative Discourse

An important requirement before a computational solution can be developed is to have a clear model of what such a solution would imply in terms of elements to represent and operations necessary to extract them from the discourse. A small outline of such a description is attempted in Section 3.1.

We want to interpret a text for a story onto a model of how an average reader might construct plausible views - at each particular point of the reading process – on the truthfulness of the various events told. This would include several stages. First it would require a process of parsing the text onto a representation of the discourse as a linear sequence of conceptual descriptions of the sentences in the text. This initial stage has been skipped in the present paper in favour of considering inputs already presented in such a format (see Section 3.3). This is done to allow the effort to focus on the development on the cognitive tasks that would be required to process such a representation through the remaining stages. Second, it would need to parse the linear sequence of conceptual descriptions of the statements in the discourse to segment it into the different narrative levels that are present in it. This involves identifying the start and end of any embedded story present, and separating the discourse into smaller segments that correspond to spans of discourse at different narrative levels (see Section 3.4). Third, it would need a process that traverses this nested structure of discourse spans for different narrative levels to reconstruct from it a partially ordered representation of the set of events described, aiming to recover an order that not necessarily corresponds to the chronological order of the events but which is compatible with the temporal relations between the events implied by their order of presentation in the discourse (see Section 3.5). This process will also uncover situations in which different parts of the discourse present conflicting descriptions of events at particular points in time. Such situations are represented in the constructed data structure as forks in the truth continuum for the story. Because each contribution to the discourse has the potential both to insert events at any point in this reconstructed ordering, and to introduce forks that add alternative views on given situations, subsequent contributions are likely to alter significantly the representation of events already described. For this reason, the final state of such a data structure would not provide a complete description of the interpretation process. Instead, such a complete description would have to include the sequence of instances of the data structure at each point in the reading of the discourse. Finally, a process is needed to traverse these reconstructed structures of the relative order of events in the story and build from them a set of views on the story that are consistent from the perspective of event truth (see Section 3.6). Again, because these sets of views will change progressively as the discourse is read, the model of the full interpretation process is not the final state of such a data structure, but rather the sequence of instances of the data structure at each point in the reading.

3.1 A Description of Truth-Sensitive Interpretation for Discourse

The task addressed in this paper involves the interpretation of a sequence of statements for the discourse of a narrative into a set of alternative possible views on a story world. In the representation of these alternative possible views the events involved must be presented in at least a correct relative chronological order with respect to the time in which they take place. A number of alternative views on the story world are possible because we are allowing the possibility for characters to tell stories that are not true. However, ther reader faced with interpreting the narrative need not know at each point whether the stories are true or not. In fact, even if they do, the reader still needs to construct a representation in which such a false story is true, because she needs it to understand how the characters that hear the false story are understanding the story world. When a character is shown deceiving an audience by telling a false story, the reader knows the story is false, but she will still build a view of the story world in which the story is true, to capture how she knows the audience is reacting.

3.2 A Corpus of Discourses for Stories

To provide a certain empirical basis for the model presented in this paper, a corpus of narrative discourse for stories is considered. To our knowledge there is no corpus of formalised stories that captures the type of features that the paper addresses. This lack is surmounted by constructing a small corpus from a reputable source that provides an analysis of the meaning of the stories in question that includes the details that are relevant for our analysis. The formal representation of the narrative structure of Russian folk stories developed by Vladimir Propp (1968) has for a long time been a preferred reference in the world of computational modelling of narrative. Propp analysed a corpus of Russian folk tales and proposed the concept of *character function* as an abstraction of certain actions by the characters that were fundamental to the evolution of the plot. Some of these functions he identified as being part of larger sequences, such as the *hero sequence* – being called to action, defeating a villain, rescuing a victim, receiving a reward –, the *donor sequence* – a donor tests a character and equips them with some magical item if their reaction to the test is positive –, or the *false hero sequence* – a character claims for himself merits that belong to someone else in the hope of appropriating the reward but is eventually exposed and punished. The appendices of

Conjunction of statements	Meaning
meets boatmen son location Xvalynsk travels boatmen at_target boatmen Xvalynsk carries boatmen son at_target son Xvalynsk	The son meets the boatment and they take him to Xvalynsk

Table 1. Example of an update involving three statements, each expressed by a predicate followed by a number of arguments, where arguments are constants referring to elements in the story world.

Propp's book provide a set of 13 tales analysed in terms of sequences of character functions. This corpus is relevant because it is the only one for which a full analysis of its narrative structure exists in terms of the character functions defined by Propp. This implies that the set of events captured by the text of the story has already been transcribed into a formal representation of the meaning, and the relations between the various events analysed in detail. We therefore choose this set of stories as case study for the proposed model. These stories are actually provided not as texts but as short-hand paraphrases of the content. So the representation for these tales as text would need to be fleshed out by hand from the short-hand paraphrases. We have opted instead for transcribing these short-hand paraphrases directly into predicate logic descriptions (see Section 3.3).

3.3 Representing Discourse

For the purposes of this paper, a *discourse* is represented as a sequence of updates to the story, where each *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. The set of statements in an update is understood to represent the conjunction of the corresponding predicates. The stories in our chosen corpus are transcribed into this representation ensuring the elements of Propp's analysis are captured correctly. Table 1 presents an example of an update.

For the representation of embedded discourses as distinct spans of the input discourse within the larger structure of the discourse we will resort to the computational machinery proposed in (Gervás, 2021b). In order to make use of the interpretation machinery presented in that paper, an embedded story will be represented as a combination of the following elements: (1) a statement to act as start of story marker (with a name for the story to act as its unique id) (`start_story <story-name>`), (2) a sequence of statements for that story (just like those for the frame story), (3) 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 end of story marker.

This notation allows the explicit representation of the content of embedded stories within the sequence of the discourse for the story. An example of the discourse sequence used as input for the interpretation of a story rendered is presented in the left hand column of Table 2. This solution for encoding the input discourse may sometimes result in a reiteration of events already told as part of the frame story – when some characters tells other characters about something that has already been mentioned in the discourse of the frame story – but it also allows for representing situations in which the embedded story differs from that told in the frame story, or to include different stories.

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 2. Example of input discourse sequence and structured interpretation for tale 155 as analysed by Propp (1968), showing how the embedded story is explicitly declared in the discourse sequence used as input for the interpretation process and how the embedded story is explicitly separated as a structure. In the discourse, the inline declaration of the embedded story, and the statement that introduces the embedded story have been boxed for clarity.

Embedded stories have two important properties that determine the role they play in the interpretation of the frame story. Each story can either be known to be true or be known to be false. This is represented by assigning a *truth value* to its representation. Certain stories are told with the intention of describing the truth of some events as known by the narrator, but others are told to present some events in a particular light, say, for describing a plan for the future, or a hypothesis about what may have happened. Such as stories are identified in the discourse representation by having an additional argument that describe their *modality*. Stories without such an intention are assigned a `told` value for modality.

In the context of a story there are sometimes contributions to the discourse – in the form of sentences about the characters or the events – that effectively cast doubt or altogether negate versions of parts of the story told previously. In the classic analysis of Russian folktales by Vladimir Propp (1968), the *false hero* sequence involves a character falsely claiming the noble deeds of the hero to steal possible rewards. This usually takes the form of the false hero telling a false version of the corresponding deeds. The sequence includes an *exposure* event which confirms that this altered version of the deeds in question was false. To be able to account for the dynamics of how truth consistency of the story is managed during interpretation, the model needs to account for such predicates. In the model, such predicates are represented by specific instances of the telling predicates, carrying reference to the told story by its name – for instance, `story_false` **`princess'_torment`** for the story shown in the example in Table 2 – whose interpretation is that the corresponding story has become known to be untrue. Rather than delete the corresponding story from the representation graph being built, the representation of that embedded story is labeled as false. This allows subsequent processes to ignore it when compiling truth-consistent views, while retaining in the model the information that the story was told at a given point in time. When a story falsifying statement is found, the representation for the corresponding embedded story is labeled as untrue regardless of what category it was stored under (the set of categories under which embedded stories are stored is described in Section 3.5).

3.4 Segmenting a Conceptual Description of Discourse into Distinct Narrative Levels

The solution proposed in this paper relies on the process for the interpretation of embedded stories as a stack-based mechanism for handling the changing contexts of interpretation presented in (Gervás, 2021b).[1] The process starts with an empty stack for the initial narrative level. On detecting the start of an embedded story (start_story <story-name>) the interpretation of the frame story to that point is pushed to the stack, and a new empty interpretation is created for the embedded story. The discourse for the embedded story is processed against this specific representation for its interpretation. When the end of the embedded story is reached (tell_story statement <narrator> <narratee> <story-name>) the accumulated interpretation for the embedded story is stored in a table for embedded sub-stories indexed by the name of the sub-story (<story-name>), and the interpretation for the frame story acting as context is popped from the stack and established as context for the rest of the frame story. An instance of the tell_story statement <narrator> <narratee> <story-name>statement is then added to the interpretation of the frame story, indicating how the embedded story fits into the frame.

This procedure produces 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 narrative-level segmented discourse is traversed, starting at the frame story, and each of the statements of the frame story is added to a new linear sequence of statements. When an embedded story is mentioned in the story being processed, it is retrieved from the structure and appended at that point to the linear sequence being built. Any embedded stories nested within that story are processed in the same way. The single linear sequence of discourse that results from this process should match the original discourse. The structure for the narrative-level segmented discourse has the advantage that it identifies explicitly both embedded stories as spans of discourse and borders between narrative levels.

An example of the sequence of stages by which a discourse is processed onto a segmented discourse is shown in the left hand column of Table 3. For ease of illustration, the story in the example has been simplified so that each update in the discourse has been reduced to a single statement (except for update 14 that carries and additional story falsifying predicate). Updates would normally each involve a set of statements representing a conjunction of predicates. The example shows how the occurrence of embedded stories related to the false hero sequence (in the middle column) progressively modifies the graph representation (on the right). The progression can be tracked by the number on the far left column, which corresponds to the different steps in the reading of the discourse. One can see that up to step 3 in the reading a single sequence of updates is being constructed, that at step 4 the first embedded story is found, then steps 5 to 9 again extend the frame story with simple updates, step 10 introduces another embedded story, steps 11 to 13 extend the discourse with simple updates, and step 14 adds a slightly more complex update – indicated with a $\wedge$ conjunction sign – that includes an additional falsifying predicate that questions the embedded story told in step 10.

1. To ensure the process described here is understandable without consulting the original paper, a brief description is included here. Interested readers are invited to consult the original paper.

Step	**Segmented discourse**	**Branching graph**
	⋮	⋮

Step 3

Segmented discourse:

Pos.	*Frame story*	*Pos.*	*Embedded story*
1	sets_out brother1		
2	sets_out brother2		
3	finds brother1 kingdom		

Branching graph:

	Preceding stories	*Alternatives*
	Frame story	*Alternatives*
1	sets_out brother1	
2	sets_out brother2	
3	finds brother1 kingdom	

Step 4

Segmented discourse:

Pos.	*Frame story*	*Pos.*	*Embedded story*
1	sets_out brother1		
2	sets_out brother2		
3	finds brother1 kingdom		
4	tells_story brother1 brother1 **abduction**	0	kidnap dragon princess

Branching graph:

	Preceding stories	*Alternatives*
	abduction	
0	kidnap dragon princess	
	Frame story	Alternatives
1	sets_out brother1	
2	sets_out brother2	
3	finds brother1 kingdom	
4	tells_story brother1 **abduction**	

⋮

Step 9

Segmented discourse:

Pos.	*Frame story*	*Pos.*	*Embedded story*
1	sets_out brother1		
2	sets_out brother2		
3	finds brother1 kingdom		
4	tells_story brother1 brother1 **abduction**	0	kidnap dragon princess
5	fight brother1 dragon		
6	defeats brother1 dragon		
7	releases brother1 princess		
8	gives princess brother1 ring		
9	arrives brother2 palace		

Branching graph:

	Preceding stories	*Alternatives*
	abduction	
0	kidnap dragon princess	
	Frame story	Alternatives
1	sets_out brother1	
2	sets_out brother2	
3	finds brother1 kingdom	
4	tells_story brother1 **abduction**	
5	fight brother1 dragon	
6	defeats brother1 dragon	
7	releases brother1 princess	
8	gives princess brother1 ring	
9	arrives brother2 palace	

Step 10

Segmented discourse:

Pos.	*Frame story*	*Pos.*	*Embedded story*
1	sets_out brother1		
2	sets_out brother2		
3	finds brother1 kingdom		
4	tells_story brother1 brother1 **abduction**	0	kidnap dragon princess
5	fight brother1 dragon		
6	defeats brother1 dragon		
7	releases brother1 princess		
8	gives princess brother1 ring		
9	arrives brother2 palace		
10	tells_story brother2 princess **false_claims**	0	fight brother2 dragon
		1	defeats brother2 dragon
		2	releases brother2 princess
11	asks_for princess ring		
12	gives brother1 princess ring		
13	recognises princess brother1		
14	exposed brother2		
	story_false **false_claims**		

Branching graph:

	Preceding stories	*Alternatives*
	abduction	
0	kidnap dragon princess	
	Frame story	*Alternatives*
1	sets_out brother1	
2	sets_out brother2	
3	finds brother1 kingdom	
4	tells_story brother1 **abduction**	
5	fight brother1 dragon	fight brother2 dragon
6	defeats brother1 dragon	defeats brother2 dragon
7	releases brother1 princess	releases brother2 princess
8	gives princess brother1 ring	
9	arrives brother2 palace	
10	tells_story brother2 princess **false_claims**	

⋮

Step 14

Segmented discourse:

Pos.	*Frame story*	*Pos.*	*Embedded story*
1	sets_out brother1		
2	sets_out brother2		
3	finds brother1 kingdom		
4	tells_story brother1 brother1 **abduction**	0	kidnap dragon princess
5	fight brother1 dragon		
6	defeats brother1 dragon		
7	releases brother1 princess		
8	gives princess brother1 ring		
9	arrives brother2 palace		
10	tells_story brother2 princess **false_claims**	0	fight brother2 dragon
		1	defeats brother2 dragon
		2	releases brother2 princess
11	asks_for princess ring		
12	gives brother1 princess1 ring		
13	recognises princess brother1		
14	exposed brother2 $\wedge$		
	story_false **false_claims**		

Branching graph:

	Preceding stories	*Alternatives*
	abduction	
0	kidnap dragon princess	
	Frame story	*Alternatives*
1	sets_out brother1	
2	sets_out brother2	
3	finds brother1 kingdom	
4	tells_story brother1 **abduction**	
5	fight brother1 dragon	*fight brother2 dragon*
6	defeats brother1 dragon	*defeats brother2 dragon*
7	releases brother1 princess	*releases brother2 princess*
8	gives princess brother1 ring	
9	arrives brother2 palace	
10	tells_story brother2 princess **false_claims**	
11	asks_for princess ring	
12	gives brother1 princess1 ring	
13	recognises princess brother1	
14	exposed brother2 $\wedge$	
	story_false **false_claims**	

Table 3. Extract from the interpretation for tale 139.

3.5 Building a Branching Partially Ordered Graph for Story Events

The *narrative-level segmented discourse* can now be processed to identify both a plausible relative ordering between the events mentioned and any conflicting view on particular sets of events. The procedure for processing the *narrative-level segmented discourse* essentially follows the traversal procedure that would be required to reconstruct the original sequence of discourse: starts at the beginning of the frame story and, on finding a reference to an embedded story, processes that before continuing – applied recursively at need. At each point in the discourse the procedure inserts a representation of the corresponding statement as an event in a *branching partially ordered graph* in one of the following categories:

- as part of a story told by the characters but not related to the frame story or its characters (*unrelated stories*)
- as part of a story about the characters in the frame story but happening before the start of the frame story (*preceding stories*)
- as part of a story which is a conflicting view on part of the frame story (*conflicting stories*)
- as part of a story about the characters in the frame story but declared to be in some way modal (*modal stories*)

Of these, unrelated stories are not relevant to considerations of truth-consistent views on the frame story been told, but they feature very significantly in literature as a whole. A famous example of this is *The Arabian Nights* (Burton et al., 2004), which is basically a simple frame story that acts as container for all the stories told within. For this reason, the system has been designed to be able to cope with them.

Modal stories are by their nature also not relevant to considerations of truth-consistent views on the frame story been told. In this case, the system has been designed to handle them because they are often present in the stories to be analysed, and they need to be identified even if just to be kept aside when compiling the truth-consistent views for the story.

The events in the *narrative-level segmented discourse* are processed into these categories according to the following heuristics:

- events from the frame story are inserted after the preceding event in the frame story
- on reaching an embedded story:
  - if the embedded story is marked as modal, it is stored separately, indexed by name (modal stories would be inserted as embedded stories in a reconstruction of the discourse at the point of telling, but they are not considered to be true in the same sense as the events in the frame story)
  - if the embedded story is marked as a told story:
    - if the span for the embedded story partially matches any segment of the preceding spans of frame story, the embedded story is inserted in the graph right before the

start of the matching span of the frame story and marked as a conflicting view on the events in the span

* if the embedded story partially matches one of the preceding stories it is inserted at the same point as that story and marked as a conflicting view on it

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.

This algorithm produces a *branching partially ordered graph* that holds all the events in the narrative. The main order in which events are presented in the graph corresponds to the order in which they appear in the discourse span for the frame story. This follows the general assumption that the events in a main story sequence appear in their chronological order (Adams, 1999). The events that appeared in embedded stories within the frame story (or recursively nested at deeper levels) are placed with respect to this reference order. Such a graph includes explicit representation of any embedded stories appearing in the narrative that do not refer to either the events or the characters in the frame story. It aso provides explicit representation of any embedded stories that constitute modal statements – not stated as being true but rather as being potentially or desirably true. It presents compiled at the start of the graph of all embedded stories that appeared later in the discourse but were intended as explanations of events that happened before the story started. Finally, it includes compiled at the point of occurrence of certain events in the frame story (possibly more than one) embedded stories that constitute conflicting views on those events.

An example of how a graph of this type is built progressively in response to the changes in the segmented discourse as they appear is shown in the right hand column of Table 3. For simplicity, the graph is shown on a template that ommits the records for types of stories that do not appear in this story – in this case, the story has no unrelated stories and no modal stories, so only preceding stories. Due to their nature, preceding stories are shown before the sequence of statements for the main story. The embedded story introduced at step 4 is identified as a preceding story – shares characters with the frame story but does not match any span of its discourse – and it is inserted before the sequence of the frame story. The embedded story introduced at step 10 is found to match the span of discourse 5 to 7 of the frame story – its three updates each share at least one statement that has the same predicate name and shares at least one argument with those of the corresponding span – so it is considered a conflicting story and it is inserted at the corresponding point. The story falsifying predicate at step 14 forces the labelling of the `false_claims` story as known to be untrue – indicated in the table by emphasizing the corresponding text.

3.6 Extracting Truth-Consistent Views of the Story from the Event Graph

The *branching partially ordered graph* produced by the preceding stage constitutes a compact representation of all the events known at a particular moment in the reading of the story, marked for

Step	Views														
3	1	2	3												
10	1	2	3	4	5	6	7	8	9	10					
	1	2	3	4	0	1	2	8	9	10					
14	1	2	3	4	5	6	7	8	9	10	11	12	13	14	

Table 4. A sequence of the set of views obtained for the example for some of the stages of processing shown in in Table 3.

truth – as known at that point of the reading – partially ordered into a plausible relative chronology and grouped into subsets of conflicting views on the same events.

From this representation it is possible to obtain the *set of truth-compatible views on events as known at a given point of reading*. Each of the views in such a set would correspond to the sequence of events found in one possible traversal of the graph, choosing at each forking point only one of the possible conflicting views. The traversal is carried out starting with the preceding stories and leaving out any modal embedded stories.

Table 4 shows examples of the views obtained for the example shown previously at relevant points in the reading. For compactness, views are shown as references to the numbered positions of the discourse of the corresponding stories as shown in Table 3. It is interesting to note that at step 3 and up to step 9 there is a single view, that at step 10 a second view appears that replaces the fight of the first brother with the dragon to rescue the princess with a hypothetical equivalent fight by the second brother – labelled from 0 because the embedded story in question is processed without context –, and at step 14 when the second brother is exposed as a false hero this second view disappears – because even though it is still present in the graph, the fact that it is labelled as being untrue makes it be ignored during view extraction. At each of these steps the views are seen to have expanded to account for the intervening contributions to the discourse.

3.7 Dynamics of Narrative Interpretation

As described at the start of the section, the complete model of the cognitive response of a reader who is processing a given story is not covered by any single *set of truth-compatible views on events as known at a given point of reading*. This is because the purpose of the story is precisely to explain how the knowledge of the characters about the world has changed over time (Adams, 1999), so it is the sequence of such model over the length of the story that constitutes the complete model of the outcome of the interpretation process.

This evolution of the understanding of the story over time is best perceived over the sequence of *sets of truth-compatible views on events as known at a given point of reading*. Table 5 shows the evolution of the number of conflicting views on the story at the different points of reading the story (referred to as steps in the previous examples). The values in the sequence show that, for the

Step	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Number of views	1	1	1	1	1	1	1	1	1	2	2	2	2	1

Table 5. Sequence of the number of different truth-consistent views at different steps of the process of reading through the discourse of the story shown in previous examples.

Type	**Totals**	**093**	**104**	**105**	**123**	**127**	**131**	**133**	**139**	**150**	**155**	**198**	**244**	**247**
Preluding	5	1							1	1	1	1		
Anaphoric	10			1	2	2		1	1		1	1		
False	3				1				1		1			
Modal	19	2	1	1	2	3		2			2	2	3	1

Table 6. Overview of types of embedded discourse present in the stories for the corpus of Russian folk tales analysed by Propp. For each of the 13 tales in the corpus, the table shows the number of embedded discourses of each different type. The second column presents the total numbers for each type. The numbering of the tales corresponds to the one used in the Appendices of (Propp, 1968).

simplest example, there is a progression from no conflict at the beginning – steps 1 to 9 – through a stage of higher conflict – steps 10 to 13 – to a resolution to a single view again at the end – step 14.

This dynamic match the expectations for basic formats of story. Any story starts with an empty set of views, because no information is available. The start of the story usually involves a simple sequence of events that all present the same view of the world, so the early stages of the reading show a single view. At some point conflicting views on the world may be introduced into the story, usually in the form of an embedded story that provides a contrasting view with some preceding span of the story. The set of views available duplicates. This process may continue, and it may lead to very complex structures – for instance, for thrillers or romantic comedies, where several hypotheses about more than one set of events in the story may be in play at any given time – see example in Section 4. Ideally, as the story progresses, some falsifying statements will start appearing, closing down some of the conflicting views.

It is easy to see how the charm of many stories – like the plot of some Hitchcock movies or modern thrillers – resides in their ability to sustain a dynamic of progressively opening and closing a number of conflicts, always ensuring that there are a few conflicts in play at any given time to keep the reader/ spectator interested. This type of structure replicates the model for the evolution of tension in drama known as Freytag's Pyramid (Freytag, 1894).

4. Discussion

It is useful to consider the relative incidence of the phenomena we are trying to address over a sample of real stories. The analysis of the 13 Russian folk tales formalised by Propp yields a number of instances of embedded discourse. The relative incidence of the different types of embedding is presented in Table 6.

These results provide some valuable insights. Of the total number of 37 instances of embedded discourse, 19 correspond to modal discourses. Such discourses provide information that is relevant for understanding the motivations of the characters, but they do not affect the perception of truth

	PS	AS	FS	MS
Present	5	10	3	19
Identified	5	5	8	19
Precision	100.0	100.0	37.5	100.0
Coverage	100.0	50.0	100.0	100.0

Table 7. Precision and coverage values for the current discourse engine prototype over the identified occurrences of embedded discourse.

over the story world. There are 5 instances of embedded discourse used to introduce preluding stories – stories that take place before the action of the frame story starts. There are 3 instances of embedded discourse used to introduce stories that are false. These correspond to instances of Propp's false hero sequence, were a character makes unfounded claims to the merits achieved by the hero. There are 10 instances of anaphoric discourse, which refer to a span of discourse of the frame story. These correspond to situations in which a certain character in the frame story reviews events already described in the frame story, usually to inform other characters that were not aware of them. Of those, 3 correspond to the exposure of the false heroes in the corresponding instances of Propp's false hero sequence. Another 4 correspond to cases where someone exposes a villainy by retelling it to someone in a position to take the villain to task. Three more correspond to situations in which a character is informed of prior events, that then lead that character to react in a certain way.

4.1 Automated Interpretation of a Corpus of Discourses

The current prototype for the described discourse interpretation engine has been run over the selected corpus. Table 7 presents the results on the ability of the discourse engine to identify and correctly represent the features known to appear in the corpus.

These results indicate that modal and preluding stories are identified by the system without any problems. False stories presented by characters as alternatives to what actually happened are also identified correctly. The problem arises in the identification of anaphoric stories that are meant to refer to spans of discourse preceding their telling in the discourse of the corresponding frame story. Half of the stories of this type that appear in the corpus are mistakenly identified as false stories. This happens due to two different reasons. First, because the current solution for identifying matches between spans of discourse is not capable of processing spans of discourse that include embedded discourses themselves. This happens in at least two instances in the corpus, and the algorithm as it stands cannot identify the match. This issue will be addressed as further work. Second, because the current solution applies a very strict algorithm for identifying anaphoric reference of this kind: it only considers exact matches of the corresponding spans of discourse. When the version told of the corresponding events does not match exactly the original telling, the system fails to recognise the reference and labels this second telling as an alternative view proposed for the same event. This happens in several cases due to the second tellings being told with less detail than the original. A simple solution would be to edit the stories in the corpus to ensure strict similarity between the orig-

	Preceding stories	*Alternatives*	
0	**murder_Serse_Artabano** kills Artabano Serse	**murder_Serse_Dario** kills Dario Serse	**murder_Serse_Arbace** kills Arbace Serse
	Frame story	*Alternatives*	
	⋮		
8	tells_story Artabano Arbace **murder_Serse_Artabano**		
	⋮		
14	tells_story Artabano Artaserse **murder_Serse_Dario**		
	⋮		
37	tells_ Megabise Artaserse **murder_Serse_Arbace**		

Table 8. Extract from the graph for the plot of Artaserse.

inal span and the retelling. However, we prefer to consider an adaptation of the matching algorithm. This problem will be addressed in future work by considering the possibility of reference being understood also in cases where the events in question are presented in a summarised or paraphrased manner.

4.2 Extension to More Complex Discourses

Although the examples shown to this point have been specifically chosen for their simplicity, to allow for a description of all the aspects involved in the process within the space limitations of a conference paper, the model can cope with stories of considerable complexity.

A parallel effort is being made to encode in the described notation the plots of the set of 18th century operas based on librettos by Metastasio, which is the focus of the Didone project.[2] An example based on one of the encoded plots is discussed below for illustration. Table 8 shows an extract – most of the first act – for the graph for the plot for "Artaserse" (Rome, 1730),[3] an 18th century *opera seria* by Metastasio that revolves around the murder of the Persian King Serse. Throughout the plot, several characters are accused of the murder. In all cases, the accusation takes the form of an embedded story. Because the actual murder itself is also told as an embedded story – a confession by the murderer at step 8, reasonably early in the plot – the machinery described accounts for this argument by making the initial confession a preceding story and alligning the other accusations – at steps 14 and 37 – as alternatives to it. This is because the corresponding embedded stories are framed as accusations rather than hypotheses – otherwise they would have been analysed as modal stories.[4]

Table 9 shows part of the evolution of the views through the corresponding fragment of the plot – this is just for the first act – for Artaserse. The number of views is not affected by the confession of the murder at step 8, because that is the first mention of Serse's death, so the murder is accommodated as a preceding story. The count of views increases at step 14 when Serse's son Dario is – unjustly as it happens – accused of the murder. At step 31 proof is found that Dario –

2. https://didone.eu/
3. http://www.progettometastasio.it/testi/ARTASERS|P1
4. A number of modal stories do occur in the encoding for "Artaserse" – plans and hypotheses – which are not shown here as not relevant to truth consistent views. The system correctly ommits them during the extraction of truth-consistent views.

Step	1	…	14	…	31	…	37
Number of views	1	…	2	…	1	…	2

Table 9. Sequence of the number of different truth-consistent views at different steps of the process of reading through the discourse of the plot for Artaserse.

who at this point has already been executed as a criminal – was not guilty, so the count of views decreases to one. At step 37 Arbace is now accused of the murder. The system appears to handle reasonably well even complex stories of this type.

4.3 Discussion of Design Decisions

The research described in the paper includes some controversial decisions, some of which are discussed here.

The decision to focus on sequences of conceptual descriptions of statements in a discourse rather than text is taken to avoid the pitfalls of having to interpret natural language discourse into a sequence of first order logic statements that may be processed automatically. This decision corresponds to focusing the paper on understanding the behaviour of discourse – as defined by Ricoeur – leaving aside the complexities of the "grammatical devices" of particular languages. It would be possible to consider existing tools for Information Extraction or other simple NLP tools applied to existing stories as a possible source of additional input material for the interpretation processes presented here. To achieve that, the existing tools would need to be extended to include the option of extracting the features specifically added to our representation to capture the essential mechanisms of discourse embedding, reference to spans of preceding discourse, and story falsification. This last point presents particular difficulties, because the task of extracting reported speech is in itself an open question (Krestel et al., 2008).

The consideration of anaphoric stories extends the two functions of language described by Ricoeur – identification and predication – to cover not only events in the story world but sequences of statements in the discourse already processed. This is because the predicates introduced to account for embedded discourses allow some arguments of the predicate to be themselves discourses. Furthermore, the predicates for indicating certain stories are false allow stories to be referred to by name, and their semantics involves a kind of broad brush negation of their overall truth.

The process of interpretation that the proposed model capture corresponds to Ricoeur's process to reduce the ambiguity of discourse Ricoeur (1980), in particular applied to ambiguity in the truth value assigned to some stories told within the discourse.

The task of constructing the graph of events associated with a particular discourse correspond to Labov's definition of narrative as a match between a verbal sequence of clauses and the sequence of events which actually occurred (Labov, 1972). It also matches partially the task of constructing worlds that Gerrig (1993) and Brunner (1986) consider fundamental among the cognitive operations involved in the task of reading fiction. The interpretation onto situation models would correspond to further processing for each of the units identified in the current paper, from basic statements characterising a possible world onto full situation models enriched with details via inference based on common knowledge.

The proposed model suggests that the outcome of the interpretation is not a single representation but a succession of views on how the reader envisaged the story world at each point of the reading. This matches Gerrig's idea that the reader produces not a single description of the world, but rather a sequence of descriptions that represent the states of the world as the story evolves. It also matches Fish's idea that the moment-by-moment of interpretation is actually an integral part of the meaning of the narrative, as opposed to being simply a process by which the meaning is obtained (Fish, 1980).

The way in which the model of the interpretation process allows more than one partition of possible views on the story world to survive in memory while interpretation is under way, with the possibility to refine the representation later by identifying some of the views as false, can be seen as a precursor of the construction of the set of different views that a reader carries out as part of the mentalisation process (Tal-Or & Cohen, 2010). Proper mentalisation would require that different views of the story wold be constructed for different characters.

The proposed representation of embedded stories accounts for their use in presenting material out of its natural chronological order, as described by Adams (1999), and it includes an acceptable solution that identifies a plausible order for the events involved that is compatible with the logical relations between them and with the way they are presented in the discourse.

The proposed model arose in the context of ongoing efforts to model literary creativity in computational terms. The computational creativity community has over a number of years attempted to create artificial systems capable of producing narratives acceptable to humans as at least valid attempts. Many approaches have been considered (Gervás, 2021a), with not very successful results. The process has produced at least some valuable insights on how the evident shortcomings of the accumulated efforts arise. One important such insight is that the human ability to create narrative is very tightly coupled with the corresponding ability to understand narrative – which most of the approaches to artificial literary creativity have systematically side-stepped. The present effort constitutes an initial attempt at developing a model that contemplates some of the elementary mechanisms that human narratives rely on. It is hoped that it will find application in support of further efforts to develop computational models of narrative construction.

The model as it stands constitutes a very simple initial approximation to the problem, so there remains a lot of work before it can be considered complete. However, some obvious improvements that could be considered for further work show significant promise. The present version of the model represents events in terms of the truth value that the reader assigns to it according to what she has read at a given point. The conflicts perceived by the reader as she works through the story very often rely on the knowledge that different characters in the story have very different views of what they perceive to be true. Therefore, an extension of the model that assigns to the different character different views on truth based on whether they have been exposed to the spans of story where they are presented would have significant potential for better modelling reader perception of conflicts in the story.

5. Conclusions

The process by which a narrative discourse produces upon the mind of the reader of a novel – or the spectator of a movie – the succession of impressions that we are accustomed to experiencing goes well beyond the ability of transcribing sentences into some representation of their meaning. Once the abstraction is made to consider that initial step as given, there remains still a significant challenge to understand how a set of conflicting views on the truth of events about the story world emerges from the progressive processing of the sequence of contributions in the discourse, how simple mechanisms such as embedded stories and story falsifying predicates are employed to drive that process of emergence, and how the resulting understanding evolves over time to create the well known impressions of setting up a situation, ramping up into a climax of conflict, and scaling down into a *denouement*.

The model proposed in this paper provides an initial approximation as to how some of these features may be modeled. The model matches the type of cognitive responses that we are used to observing in readers, responds appropriately to the mechanisms considered, and produces an output that exhibits the general characteristics that we associate with acceptable narrative structure of a story.

Acknowledgements

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 the Didone Project (`http://didone.eu`) funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme, Grant agreement No. 788986. .

References

Adams, J.-K. (1999). Order and narrative. *Recent Trends in Narratological Research*, (pp. 111–27).

Brunner, J. (1986). *Actual minds, possible worlds*. Harvard University Press. From `https://books.google.es/books?id=--\_ItAEACAAJ`.

Burton, R., Burton, R., & Byatt, A. (2004). *The Arabian Nights: Tales from a Thousand and One Nights*. Modern Library Mm. Modern Library.

Fish, S. (1980). *Is there a text in this class?: The authority of interpretive communities*. Harvard University Press. From `https://books.google.es/books?id=bYBso1t4ylcC`.

Freytag, G. (1894). *Freytag's technique of the drama: An exposition of dramatic composition and art*. Griggs.

Genette, G., Lewin, J., & Culler, J. (1983). *Narrative discourse: An essay in method*. Cornell paperbacks. Cornell University Press.

Gerrig, R. J. (1993). *Experiencing narrative worlds: On the psychological activities of reading*. Yale University Press. From `http://www.jstor.org/stable/j.ctt1dr37cb`.

Gervás, P. (2021a). Computational models of narrative creativity. In P. Machado, J. Romero, & G. Greenfield (Eds.), *Artificial intelligence and the arts: Computational creativity, artistic behavior, and tools for creatives*, chapter 9, 209–255. Cham: Springer International Publishing.

Gervás, P. (2021b). A model of interpretation of embedded stories. *Text2Story: 4th International Workshop on Narrative Extraction from Texts*. CEUR Workshop Proceedings, Lucca, Tuscany: CEUR Workshop Proceedings.

Herman, D., Jahn, M., & Ryan, M. (2010). *Routledge encyclopedia of narrative theory*. Taylor & Francis.

Kamp, H., & Reyle, U. (1993). *From discourse to logic: Introduction to model-theoretic semantics of natural language, formal logic and discourse representation theory*, volume 42 of *Studies in Linguistics and Philosophy*. Dordrecht: Springer.

Krestel, R., Bergler, S., & Witte, R. (2008). Minding the source: Automatic tagging of reported speech in newspaper articles. *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*. Marrakech, Morocco: European Language Resources Association (ELRA).

Labov, W. (1972). *Language in the inner city; studies in the black english vernacular*. University of Pennsylvania Press Philadelphia.

Lascarides, A., & Asher, N. (2007). Segmented Discourse Representation Theory: Dynamic Semantics with Discourse Structure. In H. Bunt & R. Muskens (Eds.), *Computing meaning: Volume 3*, 87–124. Kluwer Academic Publishers.

Mann, W. C., & Thompson, S. A. (1988). Rhetorical structure theory: Toward a functional theory of text organization. *Text*, *8*, 243–281.

Nakasone, A., & Ishizuka, M. (2006). SRST: A storytelling model using rhetorical relations. *Technologies for Interactive Digital Storytelling and Entertainment* (pp. 127–138). Berlin, Heidelberg: Springer Berlin Heidelberg.

Propp, V. I. (1968). *Morphology of the folktale*, volume 9. University of Texas Press.

Ricoeur, P. (1976). *Interpretation theory: Discourse and the surplus of meaning*. TCU press.

Ricoeur, P. (1980). Narrative time. *Critical inquiry*, *7*, 169–190.

Taboada, M., & Mann, W. C. (2006). Rhetorical Structure Theory: looking back and moving ahead. *Discourse Studies*, *8*, 423–459.

Tal-Or, N., & Cohen, J. (2010). Understanding audience involvement: Conceptualizing and manipulating identification and transportation. *Poetics*, *38*, 402–418.

Teasdale, B., Maguire, L., Budelmann, F., & Dunbar, R. I. M. (2021). How audiences engage with drama: Identification, attribution and moral approval. *Frontiers in Psychology*, *12*.

Van Dijk, T. A., & Kintsch, W. (1983). *Strategies of discourse comprehension*. New York: Academic Press.