Targeted Storyfying: Creating Stories About Particular Events

Pablo Gervás
Facultad de Informática
Instituto de Tecnología del Conocimiento
Universidad Complutense de Madrid
pgervas@ucm.es

Abstract

The present paper proposes a computational model of the task of building a story from a set of events that have been observed in the world. For the purposes of the paper, a *story* is considered to be a particular type of sequential discourse, that includes a beginning, a complication and a resolution, concerns a character that can be clearly identified as a protagonist, and ends with a certain sense of closure. Starting from prior approaches to this task, the paper addresses the problem of how to target particular events to act as the core of the desired story. Two different heuristics – imaginative interpretation and imaginative enrichment – are proposed, one favouring faithful rendering of the observed events and the other favouring strong cohesive plots. The heuristics are tested over a simple case study based on finding interesting plots to tell inspired by the movements of pieces in a chess game.

Introduction

Interest in stories as means of communication has risen over the past decade. As a result research efforts targeting the computational treatment of stories have proliferated. The academic world is beginning to realise that there is much more to how we use stories in communication than just understanding them or generating them. People interact with stories in many different ways, and in fact most of the situations in which we create stories in our everyday life do not involve inventing the events that make up the story. Rather, they involve constructing a story out of events that we are aware of, in order to communicate either the events themselves or some particular interpretation of them. Yet sometimes, for the sake of better communication, we do invent some events to give the story cohesion, or charm. This task has been overlooked in computational approaches to stories, in favour of more creative approaches to storytelling, where the complete story – including its constituent events – is made up from scratch. Yet there is special kind of creativity involved in coming up with a story to match a given set of facts. As we know, finding "the right story" to tell about a set of facts can be crucial for successful communication, making the difference between flat rendition of the facts and either an entertaining yarn or a convincing argument. The development of computational models of this task would be an important contribution to the field of computational narratology.

Although the task of composing a narrative based on a given set of events that have taken place is very relevant to understand how humans structure their perception and memory of the world around them, it has received less attention in terms of computational modelling than other tasks related to stories (story understanding (Mueller 2003) or story generation (Gervás 2009; Kybartas and Bidarra 2016)).

Part of the problem is that the study of this task from a computational point of view requires a model of the input (the set of facts observed / remembered that constitute the source and the starting point for the composition). Such a representation of the input is implicit in the specification of the task, yet any attempt at computational modelling must start by representing it explicitly, as it will significantly influence the rest of the process.

Advances have been made in the understanding of the task by considering a chess game as a very simple model of a formalised set of events susceptible of story-like interpretations. Chess provides a finite set of characters (pieces), a schematical representation of space (the board) and time (progressive turns), and a very restricted set of possible actions. In this approach, each individual chess piece taking part in the game is considered a character. Perception range is defined as the small space of N x N squares of the boad that constitutes that immediate surroundings of each piece at any given moment. Events are triggered by pieces moves. Whenever a piece moves, this constitutes an event for the piece itself, for any other piece captured during the move, and for any other piece that sees either the full move, the start of the move or the conclusion of the move. Fibres for each of the pieces are built by collecting event descriptions for those moves that they are involved in or they see. The same event may get described differently in different fibres depending on the extent to which the corresponding focalizer is involved in it.

The use of chess game data as a valid test domain relies on the assumption that these games are being interpreted as summaries of the movements and interactions of people over a map of space. It is important to note that intuitions arising from the rules of chess or chess playing experience must be disregarded for the approach to be successful. If this is achieved, the mechanics of narrative composition developed for this type of example may provide a valid source for extrapolation to more complex domains.

Related Work

The work in this paper is informed by existing theories of narrative, and some fundamental concepts of narratology need to be taken into account. The paper addresses the task of storyfication, which involves building a narrative construct from a set of events that have been observed in the world but do not necessarily fit together as a story before being processed. It is related to the task of narrative composition, which involves elaborating a set of observed events into a sequential discourse to be told. Although the difference is subtle, we assume that the result of a process of narrative composition is a sequential discourse that conveys a set of events, and the result of a process of storyfication is a particular type of sequential discourse, that includes a beginning, a complication and a resolution, concerns a character that can be clearly identified as a protagonist, and ends with a certain sense of closure.

Narrative

Narrative has been classed among the elementary cognitive abilities exhibited by human beings (Schank and Abelson 1977; Bruner 1991; Herman 2004). In particular, it is known to be the process by which humans transform a particular experience of reality into a cognitive form that is easy for the mind to store and to be communicated to other people. Based on these ideas, recent years have seen a significant effort to relate narrative to the study of human cognition (Herman 2003; 2013). An important obstacle that faces this challenge is the fact that humans are notoriously poor at identifying the processes that they apply in processing reality (Nisbett and Wilson 1977). The underlying latent processes have to be postulated from the observation of their external manifestations, such as the actual narratives as literary works –studied by narratology – or the processes by which humans produce narratives – studied by cognitive science.

Relevant concepts from the field of narratology (Abbott 2008) are the distinction between *fabula* – the set of events behind a story, independently of how they are rendered – and *discourse* – the particular way of rendering a given fabula as a sequence of statements –, and *focalization* (Genette 1980) – the way in which a story is told from the view point of particular characters, switching between them at need to cover what happens elsewhere.

Existing narratives can very rarely be paired with alternative records of the experience that led to them, or even the events that are represented in them. This is a significant obstacle for applying a data-driven approach to model narrative construction computationally, as these approaches require instances of both the input that lead to the communication impulse, the narrative that arose from it, and possibly representations of intermediate design decisions.

Cognitive scientists have proposed models of the *writing task*. Flower and Hayes define a cognitive model of writing (Flower and Hayes 1981) in terms of three basic process: planning, translating these ideas into text, and reviewing the result with a view to improving it. These three processes are framed by "the rhetorical problem" – the rhetorical situation, the audience and the writer's goals. The target events considered in the present paper would be an instance of part of this problem.

Narrative Composition

Operating on simple representations of a chess game in algebraic notation, exploratory solutions for the tasks of content selection and content planning are explored based on a fitness function that aims to reflect some of the qualities that humans may value on a discourse representation of a story. Based on this approach prior work has been carried out on exploring computational models of the task of narrative composition as a set of operations that need to be carried out to obtain a span of narrative text from a set of events that inspire the narration (Gervás 2012; 2013; 2014).

Work has also been carried out on the composition of narrative discourse from generated plots represented as plans (Winer and Young 2016). Although such efforts are not grounded on a set of events that actually happened, their approach resembles the work presented here in that the planning stage that creates a plot involves selecting a subset of all possible events based on how they might be connected (in this case, via causality), and subsequent processes determine how the selected events are organised into a discourse.

Storyfying

A computational model of the task of storyfying has been proposed in the *StoryFire* application (Gervás 2018). This model is based on a series of stages:

1. establishing how the events are perceived from the point of view of the participating agents, by partitioning experience into narrative threads centred on particular characters (a task known in narratology as *focalisation* (Genette 1980))
2. representing the structure of the story (or plot) to be constructed as an abstract frame to which the perceived events must be matched
3. mapping the events in (possibly a select part of) the narrative thread for some character into an abstract frame for a plot
4. generating a readable version of the resulting discourse

The *StoryFire* application relies on the solution for focalisation presented in (Gervás 2012; 2014), which partitions the perception of the world by a given agent into a *fibre* constructed as a sequence of events descriptions. An *event description* consisting of a set of predicates that encode the elements that appear within the perception range of the agent at a given point in space and a given moment in time. An example of an event description is given in Table 1, showing what the left white rook (`lwr`) sees around itself (the first, second and third white pawns, the left white knight, and the left white bishop) on the seventh move of the game from position `a 1` and what it sees happening (the third white

```
Focalizer: lwr
Position: a 1
Time: 7
Perception Range: 2

DESCRIPTIVE:
  is_at(wp1, a2)
  is_at(wp2, b2)
  is_at(wp3, c2)
  is_at(lwk, b1)
  is_at(lwb, c1)

NARRATIVE:
  leaves_from(wp3, c2)
```

Table 1: Example of event description, which acts as the basic unit of description for a narrative fibre.

```
PLOT ELEMENT NAME = CoupleWantsToMarry
ROLE-DATA
lover hero
beloved love-interest
```

Table 2: Description of the CoupleWantsToMarry plot element.

pawn moves two squares forward, and thereby disappears from view).

That earlier version of the *StoryFire* application relied on a representation of plot in terms of *plot frames*, which are representations as sequences of character-function-like elements (Propp 1928) known as *plot elements*. Each plot element holds a label (such as `CoupleWantsToMarry`) and a mapping between roles relevant to the plot element (such as `lover` and `beloved`) and roles relevant to the plot in general (such as `hero` and `love-interest`). An example of plot element is shown in Table 2.

The plot frames considered for earlier version of the *StoryFire* application were instantiations of the seven basic plots defined by Booker (Booker 2004).

The actual storyfication process produced a *match* between a thread and a plot frame involves an allignment between a subset of the events in a thread and the sequence of plot elements in a plot frame (described in terms of which time points in the thread are alligned with which plot elements in the frame), a mapping between the characters present in the thread and the plot roles in the plot frame and a score that corresponds to the percentage of satisfaction of set of roles involved in the plot element by roles assigned to the characters present in the matched event, averaged over all the allignment. An example of such a mapping is given in Table 3.

The preceding version of the StoryFire application produced stories that could be considered narrative interpretations of particular threads from a chess game. Given a particular piece playing in the game, the application would produce a story plot that involved that piece as protagonist and which was actually a selection out of the set of events in the narrative thread experienced by that piece during the game.

This approach was sufficient for emulating the simpler kind of storytelling that people apply, for instance, on returning from a trip. A story to tell, extracted from the events experienced during the trip, is sufficient. However, the present paper attempts to address a refined version of the task, which involves not just finding a story to tell about the trip, but finding a story that includes particular events that happened during the trip.

```
Thread lwk
PlotFrame Comedy-UnrelentingGuardian
Score 83

ALLIGNMENT
 9 [0]
11 [1]
16 [2]
17 [3]

MAPPING
bp4=love-interest
rwb=obstacle
lwk=hero
```

Table 3: Match between thread and plot frame

In addition, the set of plots considered in the earlier version was built of plots that were structurally very similar to one another. This restricted the sequences of events for which matches could be found.

Targeted Storyfying

Prior approaches to the task of narrative composition assumed that the goal was to obtain the best possible story for a given set of events. In the present paper, we want to narrow the focus to obtain stories that include a specific subset of events. In terms of the example used above, rather than build the best possible story out of the trip to a given conference, we want a story about the trip that involves, say, the keynote presentation at the conference, even if the trip might yield better stories by focusing on the conference dinner instead. This ability to drive the storyfication process towards particular events would bring the functionality being developed a step closer to human capabilities.

To this end, we need to address three different issues. First, we need to establish some means for specifying which events are to be considered as strictly required. This specification should be considered as an input in addition to the wider set of events to be considered. Second, the additional restriction imposed on the procedure may rule out matches with certain plots, and there is a risk that no match be found by applying the set of plots and the baseline algorithm previously available. Third, a procedure for guaranteeing that the produced stories include the desired events.

Establishing a Target Seed for the Storyfication

The type of constraint that humans considered when carrying out (the process that closely resembles what we are now calling) targeted storyfication is very broad. For instance, one may desire a story about a particular event, but may want the event to initiate the story, or to conclude it, or appear somewhere in the middle of it. Or one may want the story to take place at a particular location, or involve a given object. For the purpose of the present paper, we want to identify the simplest possible specification of these constraints that is compatible with the representation we are considering, and which satisfies the requirement of driving the process towards a particular subset of the input material.

The narratives we are considering are already focalised on particular characters. The simplest additional restriction that can be imposed is to consider as target a particular moment in time. This is consistent with the representation for a chess game, which is partitioned into a sequence of time points corresponding to alternating piece movements between black and white. It also allows isolation of particular events in terms of piece movements. Finally, given that the perception of the game is focalised on a particular character, specifying a moment in time also restricts the location to wherever the focaliser is at that moment.

We will therefore specify the target for our specification as a list of time points in the game. For simplicity, we will consider these time points in chronological order. This input we will refer to as the *target seed*, as the story to be built ought to be constructed around it.

Increasing the Range and Complexity of Possible Stories

The set of plots that had been considered in earlier attempts was grounded on existing accounts of plot, and it was reasonably varied in terms of the set of plot elements that it included, but proved to be ill suited to the task. First, because it included a number of classical plot structures that required that the hero travel away from home and then return. Such a structure leads to great stories, but it is very unlikely to occur in the context of a chess game. Second, because part of the plot elements included to add variation corresponded to Propp's Donor cycle, where the hero meets a character that gives to him a magical object which can then be used to solve difficulties later in the story. Again, in the context of a chess game, transfer of objects between characters (pieces) is not contemplated.

Solving these two problems was easy simply by eliminating from the set of plots those that involved journeys or donors. But as a result the set of plots was significantly impoverished, both in terms of number of plots and variety of plot element combinations. To address this problem, a refinement on the representation of plot was introduced.

The plots considered for the present paper are represented in terms of plot spans. A *plot span* represents a span of plot, constituted by a sequence of plot element (or smaller spans). The idea is to capture the concept of a number of plot elements appearing as a structural unit in a plot, but not necessarily occuring contiguously in the discourse for the plot. For example, a plot span representing an Abduction as it features in classic stories would include the actual kidnapping (which would happen somewhere towards the start of the story) and the corresponding Release (which would happen somewhere towards the end of the story), but these two plot element are structurally connected. Such cases we refer to as an *axis of interest*. Axes of interest can be combined together, weaving their corresponding subspans with those of other axes of interest, to form complex plots (which are themselves represented as plot spans). The set of plot structures described in the literature (Gervás, León, and Méndez 2015) can be represented with the help of these elements.

```
AXISofINTEREST = Abduction
PROTAGONIST = abducted
ROLES = abducted abductor rescuer

PLOT-SPAN-NAME = Kidnapping

PLOT ELEMENT NAME = Abduction
ROLE-DATA
abductor x
abducted y

PLOT-SPAN-NAME = Rescue

PLOT ELEMENT NAME = Rescue
ROLE-DATA
abducted y
rescuer z
```

Table 4: The Axis of Interest for Abduction

```
PLOT-SCHEMA = OCM-Abd
PROTAGONIST = hero

Abduction Kidnapping  (abductor=villain,abducted=victim)
CallToActionReward Call (called=hero,caller=sender)

Abduction Rescue (abducted=victim,rescuer=hero)
CallToActionReward Reward (rewarded=hero)
```

Table 5: Example of plot schema for a basic Abduction plot

An example of axis of interest is shown in Table 4. To assist in the process of combining them into more elaborate structures, each axis of interest specifies which character is the protagonist and what the roles relevant to the axis of interest are.

Axes of interest are combined into plots by means of plot schemas. A *plot schema* encodes the way in which several axes of interest combine together to form the plot span for an elaborate plot. An example of plot schema is presented in Table 5.

This shows how the Abduction and CallToActionReward axes of interest are interleaved to form the basic plot, and how the narrative roles for the plot (hero, villain, victim, sender) are mapped to the roles specific to the constituent plot elements (abductor, abducted, called, caller, rescuer, rewarded). This information is necessary to ensure that, once characters extracted from the observed set of events are mapped onto to the set of narrative roles for the plot, coherent instantiation of the plot elements with the given characters can be carried out.

The simplicity of schemas allows for the rapid construction of a large number of variations of simple plots by combining a reduced set of axes of interest, while allowing for significant structural complexity in the resulting plots, arising from the interleaving of the axes of interest.

Creating Stories for Particular Fragments of a Chess Game

The procedure to be applied for targeted storyfication is an extension of the procedure applied in the StoryFire application as reported in (Gervás 2018). That procedure involved traversing the search space of pairings between the given thread and each of the candidate plot structures. Although these plot structures are now internally represented as plot spans rather than plot frames (see sections above on *Storifying* and *Increasing the Range and Complexity of Possible Stories* for details on the differences), at the time of computating the allignments they are still converted into a sequence of plot elements to simplify the computation.[1]

For each pairing between a thread and a plot, the procedure extracts the set of characters in the thread and the set of narrative roles to be filled in the plot, and considers all possible mappings between these two sets. For each such mapping, the procedure heuristically explores possible allignments between a subset of the event descriptions in the thread and the plot elements in the plot. Any such allignment respects the relative order of events in the thread and plot elements in the plot and provides a correspondence between some of the events in the thread and each of the plot elements in the plot. Each pairing between and event description and a plot element is scored in terms of the percentage of satisfaction of set of roles involved in the plot element actually assigned to the characters present in the event. The allignments themselves are scored in terms of the average of the scores for the pairing made for all their plot elements.

For each candidate plot, only the best scoring allignment is considered, under the assumption that once a good story has been constructed out of a given thread, additional stories from that same thread – with the same plot but a different cast of characters and/or differnet allignment with the events in the thread, and with lower scores – are not desirable.

To implement the targeting of a specific subset of the input thread, two possible heuristics are applied. The first heuristic attempts to emulate the behaviour of a person trying to tell a story about some events in her day, but committed to being strictly truthful about it. It involves search for a story constructed entirely out of events that did happen. We refer to this approach as *imaginative interpretation*, because it is built up entirely of events in the original thread, but each event is interpreted as a plot element (which may involve attributing certain actions to the characters that were not explicit in the event, and attributing motivations to character behaviour). The second heuristics attempt to emulate to tell a story about some events in her day, intending to report faithfully the inspiring events but not necesarily the rest of her day. The person bases the story as faithfully as possible on the inspiring events, brings in some additional real events to support it, but may consider some fictional events to better match the story with a target plot. We refer to this approach as *imaginative enrichment*, because the material from the original thread can be enriched to adapt better to the plot of the story. In this approach, any additional events from the thread brought in to support the story need not be too faithfully rendered.

Overall, it seems that the two procedures proposed model different approaches to the task of generating stories about established facts. If the speaker wants to be careful in representing the events as they really happened, the imaginative interpretation procedure would be preferred. But the set of resulting stories may not include elaborate plots or fancy flights of fantasy. If the speaker is not so careful about representing the events as they happened, the imaginative enrichment might be preferred.

The *imaginative interpretation* heuristic involves applying the original procedure to the complete thread, but rejecting any allignments that do not include the events present in the target seed. This approach is a simple application of the original procedure with the additional constraint. It corresponds to ensuring that the part of the input thread that is being mapped onto a plot includes the desired events in the target seed.

An example of a plot produced by a process of imaginative interpretation is shown in Table 6.

The *imaginative enrichment* heuristic goes one step further. The heuristic followed here involves two different improvements. First, the process of allignment is modified so that the matching of the events in the seed to plot elements in the plot can be optimised. This involves giving preference to allignments in which the targeted events match the plot elements assigned to them perfectly, even if the score of the complete allignment over the plot drops somewhat as a result. The process of allignment is therefore broken down into an initial allignment between the targed seed and the plot (which returns an allignment of the complete seed with a subspan of the plot), and a later stage of finding allignments for any remaining subspans of the plot with the subspans of the thread surrounding the target event. Second, in cases where part of the input thread – the events in the target seed and some additional support events – is mapped to partially instantiate a plot but some plot elements in the plot have not found a correspondence with real life events, the imaginative enrichment heuristic accepts the match as valid. Because of the nature of the process, the narrative roles in the unalligned plot elements will at that stage have been mapped to characters in the thread. The resulting plot is therefore coherent, even if some of the plot elements involved are not actually supported by events in the input thread.

Plots produced by imaginative enrichment will only differ from those produced by imaginative interpretation in that, in some case, no supporting events for some for their plot elements can be shown. The surface form of plots produced by either procedure is indistiguishable.

[1] The internal structure of a plot in terms of the interwoven plot spans from different axes of interest may play a significant role in the storyfication process once subplots start to be considered. For the present, it can safely be disregarded without affecting the outcomes.

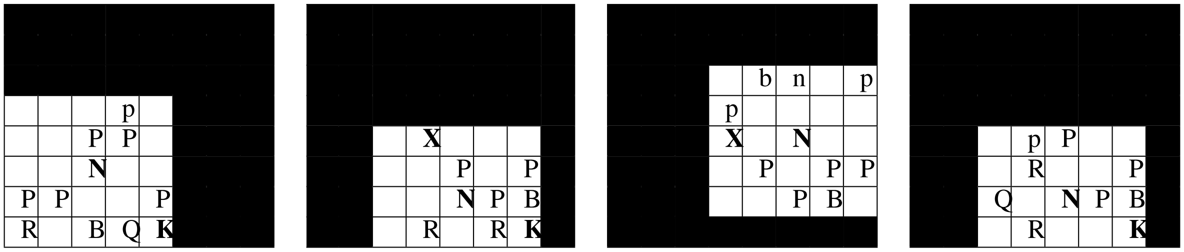

Move: 7	Move: 37	Move: 39	Move: 55
character lwk (**N**) character wk (**K**) mutual_love lwk wk want_to_marry lwk wk lover lwk beloved wk	character rwk (**X**) (guardian rwk wk) opposed_to_plan rwk sundered lwk wk	(different_class lwk wk) high_status_revealed lwk ¬ sundered lwk wk	marry lwk wk

Table 6: Storyfication of the thread for the left white knight (lwk), targeted on events 37 and 55: the left white knight (lwk, represented in the diagrams as **N**) in terms of his romance with the white king (wk, represented in the diagrams as **K**) in the face of opposition of is guardian the right white knight (rwk, represented in the diagrams as **X** to distinguish him from the left white knight).

Selecting Adequate Input Parameters

The two proposed procedures have been tested using as inputs descriptions of chess games in algebraic notation. Because every game is different in terms of what the pieces do over the duration of the game, results obtained from a particular game have to be judged in the context of how suitable that game was for the production of stories based on the available set of plots. The procedures to be tested also require as input a choice of which piece to focalise on. Again, different choices of focaliser may influence the results, as some pieces may be more active than others, or have a chance to see more activity around them. Finally, the choice of which events are included in the target seed also affects the results, as, for each thread, certain events are more likely to lead to good stories than others. For these reasons, comparisons in terms of quantitative evaluation are only meaningful across different results for a given choice of inputs.

Nevertheless, the choice of input parameters to test should be informed. Table 7 reports on the overall results obtained by the non-targeted version of the system for a sample of 6 games, to show how this property fluctuates from game to game. Results include number of stories obtained overall for each game, and average number of stories obtained for each type of piece. The length of the game is included as it clearly affects the number of stories that can be obtained from it.

These results show that longer games tend to provide more stories, and that bishops, knights and rooks are likely to produce more stories than other pieces. The choices of game parameters and focalisers for the test reported in this paper are made accordingly.

In order to provide some indication of how suitable the particular context is for each test of the proposed procedures, the prior version of *StoryFire* is applied as to the same context in each case and the average score of the stories produced is reported.

Metrics for Targeted Storyfication

The two proposed heuristics differ in nature, so they have to be evaluated differently.

Imaginative interpretation finds solutions in which every plot element in the chosen plot is alligned with an event in the input thread. Results of this approach may be evaluated on the basis of the metric already defined for scoring results of storyfication (average percentage of set of roles involved in the plot elements actually assigned to the characters present in the events).

Imaginative enrichment allows the construction of stories in which events other than those in the target seed need not be supported by events in the input thread. Results of this approach may be evaluated on the degree to which the plot of the resulting story is supported by events in the input thread. This is represented as a simple ratio between the plot elements that are supported and the total number of elements in the plot.

In order to ascertain specifically whether a solution for targeted storyfication is sucessful, the following quantitative metrics have to be considered:

- average percentage of match in character assignment between thread events and plot elements that have been alligned to one another (%MTP for match between thread and plot)
- percentage of the available plots that it has been possible to succesfully instantiate by applying a given heuristic (%PSI for plots succesfully instantiated)
- percentage of the plot elements in the plot that have been alligned with events in the input thread (%DTS for degree of thread support)

Game	Stories	Moves in game	P	B	K	R	Q	K
1	139	40	7.5	11.0	12.0	9.5	5.0	5.0
2	167	58	7.9	15.0	10.5	10.0	6.5	6.5
3	213	104	9.4	17.0	15.5	14.5	7.5	7.5
4	226	120	10.3	17.0	15.5	17.5	10.0	10.0
5	209	90	9.5	17.0	17.0	16.0	9.5	9.5
6	119	32	6.5	10.0	9.0	6.5	5.0	5.0
Av.	178.8	74	8.5	14.5	13.3	12.3	7.3	7.0

Table 7: Fluctuation of number of valid stories (overall and per type of focaliser piece) obtainable from 6 different chess games for the set of available plots. Piece types are shown as: P for pawns, B for bishops, K for knights, R for rooks, Q for queens and K for kings.

Table 8 reports values for these metrics obtained for a number of possible configurations of the input parameters. Based on the discussion presented in the Section above on *Selecting Adequate Input Parameters*, the longest game in the test set was selected, and the thread for one of the knights. The set of possible target seeds is restricted to time points that are covered by a particular thread, because the thread does not include information for time points in which the focaliser sees nothing happening around its position. A number of possible target seeds has been picked at random, both for single and double events, and considering different types of groupings of the target events with respect to the overall length of the thread.

Discussion

The results for imaginative interpretation indicate a progressive decrease in the scores as more constraining targets are provided. Whereas the un-targeted storyfication reaches an average score on MTP of 97.9, the best score for a single event seed is 96.8 and the best score for a double event seed is 93.0. The relative position of the targeted events within the overall thread also seems to affect the scores. When a targeted event is close to the beginning or the end of the thread, this imposes a limit on the number of events that can be matched to parts of a plot before or after the target. As a result, the set of possible solutions is reduced. This has a drastical effect on the number of solutions produced, as indicated by the values for the %PSI metric for the double event target seeds as the targeted event approach the end of the thread.

The results for imaginative enrichment are more difficult to interpret. The procedure employed gives priority to focusing the plot on the event in the thread that best matches the character assignments in the plot, in contrast to the imaginative interpretation procedure that gives priority to a better score over the complete plot. This priority sometimes leads the results to plot with very poor support from the input thread.

The procedures described are at present simple baselines. Many computational aspects are in need of improvement. The chess domain is in itself also a very elementary case study. It is surprising how such a simple set up can yield insight on the mechanics of putting together the elements that go into a simple story, and how it allows consideration of issues relevant to the task such as the concept of targeting particular events during plot construction, or the difference between prioritising faithful reporting of observed events or construction of rich plots loosely based on some particular selection of the observed events.

The set of plots currently in use is also a first approximation and would also need to be expanded. The proposed representation in terms of axes of interest and plot schemas, articulated as plot spans, has proven to be a powerful tool for efficiently generating a variety of plots.

In its present form, the proposed approach to storyfication is based on the application of a set of pre-compiled plots to find the one that best matches an input set of events. The use of the words “story” and “plot” should not mislead the reader into thinking that the model described in this paper is intended as a plausible model of how humans address the task of giving birth to works of narrative of literary value. The processes involved in that nobler task are undoubtedly much more complex than the procedures outlined here. The use of “story” and “plot” is made to clarify the need for the desired outputs to satisfy basic restrictions in terms of being a particular type of sequential discourse, that includes a beginning, a complication and a resolution, concerns a character that can be clearly identified as a protagonist, and ends with a certain sense of closure. The ability that these procedures attempt to model is the simpler task of packaging a subset of the events one has observed over a given period, in such a way as to tell it in an entertaining manner to someone else. These simpler stories share with their literary counterparts some of their basic constraints but none of the complexity or the elaboration. In this sense, they constitute a good case study on which to break new ground over a simple representation.

Conclusions

The decision of basing a story on a particular set of events that will appear in its core has been shown to impose significant constraints on the task. Simple exploration of the alternatives available from a computational point of view indicates that an author faced with this task would have to choose whether to aim for faithful rendering of the context in which the selected events happened, or to give priority to the events themselves, and accept the possibility of building a new context for them that improve their potential as a story.

The proposed procedures and resources are intended as a first approximation to the task. Several avenues for future research have been uncovered. More elaborate refinement

	Seed	ImagInt		ImagEnr	
		%MTP	%PSI	%DTS	%PSI
No seed	-	97.9	100		100
Single event seed	(7)	95.9	100	98.3	100
	(37)	96.8	100	27.0	100
	(72)	69.4	100	44.4	100
Double event seed	(7, 14)	93.0	100	85.5	100
	(37, 46)	85.4	90	100.0	100
	(55, 72)	80.5	60	97.7	100
	(7, 72)	73.0	20	75.3	100

Table 8: Results of storyfication: for the left white knight, on game 4 involving 120 moves, showing different choices for the target seed and their impact on the metrics: MTP for match in character assignment between thread events and plot elements, PSI for percentage of plots succesfully instantiated, DTS for degree in which the plot elements in the plot is supported by events in the thread. In each case, averages over the complete set of stories produced for the given input are given. All scores are normalised to 100 for ease of comparison.

of the procedures from a computational point of view will be addressed. The set of plots considered will be expanded. Applications beyond the chess domain case study will be considered.

Acknowledgments

This paper has been partially funded by the project IDiLyCo: Digital Inclusion, Language and Communication, Grant. No. TIN2015-66655-R (MINECO/FEDER) and the FEI-EU-17-23 project InViTAR-IA: Infraestructuras para la Visibilización, Integración y Transferencia de Aplicaciones y Resultados de Inteligencia Artificial.

References

Abbott, H. 2008. *The Cambridge Introduction to Narrative*. Cambridge Introductions to Literature. Cambridge University Press.

Booker, C. 2004. *The Seven Basic Plots: Why We Tell Stories*. The Seven Basic Plots: Why We Tell Stories. Continuum.

Bruner, J. 1991. The narrative construction of reality. *Critical inquiry* 1–21.

Flower, L., and Hayes, J. 1981. A cognitive process theory of writing. *College Composition and Communication* 32(4):365–387.

Genette, G. 1980. *Narrative discourse : an essay in method*. Cornell University Press.

Gervás, P.; León, C.; and Méndez, G. 2015. Schemas for narrative generation mined from existing descriptions of plot. In *Computational Models of Narrative*. Atlanta, Georgia, USA: Scholoss Dagstuhl OpenAccess Series in Informatics (OASIcs).

Gervás, P. 2009. Computational approaches to storytelling and creativity. *AI Magazine* 30(3):49–62.

Gervás, P. 2012. From the fleece of fact to narrative yarns: a computational model of composition. In *Workshop on Computational Models of Narrative, 2012 Language Resources and Evaluation Conference (LREC'2012)*.

Gervás, P. 2013. Stories from games: Content and focalization selection in narrative composition. In *I Spanish Symposium on Entertainment Computing*.

Gervás, P. 2014. Composing narrative discourse for stories of many characters: a case study over a chess game. *Literary and Linguistic Computing* 29(4).

Gervás, P. 2018. Storyfying Observed Events: Could I Dress This Up as a Story? In *AISB Symposium on Computational Creativity (submitted)*.

Herman, D. 2003. *Narrative Theory and the Cognitive Sciences*. CSLI Publications. CSLI Publications.

Herman, D. 2004. *Story Logic: Problems and Possibilities of Narrative*. Frontiers of narrative. University of Nebraska Press.

Herman, D. 2013. *Storytelling and the Sciences of Mind*. Cambridge, MA.

Kybartas, B., and Bidarra, R. 2016. A survey on story generation techniques for authoring computational narratives. *IEEE Transactions on Computational Intelligence and Artificial Intelligence in Games*.

Mueller, E. T. 2003. Story understanding through multi-representation model construction. In *Proc. of the HLT-NAACL 2003 Workshop on Text Meaning - Volume 9*, 46–53. Stroudsburg, PA, USA: ACL.

Nisbett, R. E., and Wilson, T. 1977. Telling More Than We Can Know: Verbal Reports on Mental Processes. *Psychological Review* 84(3):231–259.

Propp, V. 1928. *Morphology of the Folk Tale*. Akademija, Leningrad.

Schank, R., and Abelson, R. 1977. *Scripts, Plans, Goals and Understanding: an Inquiry into Human Knowledge Structures*. Hillsdale, NJ: L. Erlbaum.

Winer, D., and Young, R. M. 2016. Discourse-driven narrative generation with bipartite planning. In *Proceedings of the 9th International Natural Language Generation conference*, 11–20. Association for Computational Linguistics.