Generating a Search Space of Acceptable Narrative Plots

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

Abstract

Research on narrative generation needs to consider that a run-of-the-mill would-be novelist is usually much more concerned with getting the form of his output right than with finding new techniques or new materials. Resorting to prior techniques and materials is considered an acceptable practice, and rewrites of classic stories, sequels or series of novels set in an already described world are highly valued. In this light the need to produce outputs that are recognisable as instances of stories seems to have priority over other criteria for a narrative generation solution. The present paper follows accepted engineering practice in choosing this particular challenge as starting design goal for an initial module, to be later extended with solutions for selecting or achieving "good" or "novel" stories at a later stage. The paper proposes a representation of plot that captures both a surface structure in terms of adjacency in a discourse sequence and conceptual connections between elements of the plot that may span across its length. Based on this representation, a solution is proposed for plot generation that produces a broad range of outputs that are acceptable as instances of stories.

Introduction

Among the disciplines related to artistic creation, narrative stands out in its relative tolerance for less-than-novel artefacts – as in rewrites, sequels or series – as long as they are well-crafted products that follow established rules of how they are constructed. This is in high contrast with modern art or contemporary music, where vague resemblance to anything that went before is considered very detrimental. This dychotomy is related to the conception of creativity as a balance between novelty and value.

Even though there is not clear agreement on how to define creativity, most accounts consider that for an output to be creative it must show indications of being new and being valuable. This duality was formalised by Ritchie (Ritchie 2007) as a trade-off between typicality and novelty. To be acceptable as valid output, a particular artifact must satisfy criteria to be recognised as an instance of the target class. In Ritchie's terms, these criteria define typicality for such artifacts. To avoid confusion with biased uses of the term, we redefine this as *acceptability* of an artifact as an instance of the target class. To be considered creative, artifacts need to be different from outputs already known (novel). There is a tension between these two characteristics. If an item follows very closely the rules for a typical instance of its kind, it is unlikely to be considered novel. If an artifact innovates radically with respect to prior instances, it runs the risk of not being considered typical, even to the point where it is no longer considered an instance of the target class.

The criteria applied to make these decisions vary across genres. In modern art, for instance, novelty is valued very highly, and typicality of any kind is almost frowned upon. An art piece remotely resembling prior production in any aspect immediately loses points in the appreciation scale being applied tacitly. This includes physical resemblance, choice of material, or techniques employed in its construction. The objection that an artist in search of recognition fears the most is that of not being novel enough. In the realm of narrative, in contrast, although there is still significant pressure on finding new techniques or new materials, there is a major concern that the results be intelligible, and these imposes important concerns on how much novelty can be introduced without compromising the ability to communicate with the reader.[1] It is also true that the reuse of elements from prior stories is considered an acceptable practice (Tedford Jones 2002). This includes reuse of the structure of a previous narrative – rewrites of classic stories –, or the characters – sequels –, or the setting – series of novels set in the same world. A would-be novelist is much more concerned with getting the form of his output right (having learnt how to write in the acceptable fashion) than with minimizing his reference to the classics. In fact, having frequent reference to prior work, or resorting to known successful tropes is often done consciously in search of this impression of having mastered the craft.

From the point of view of research on narrative generation, it is clear that for this particular genre acceptability is valued highly and novelty is assigned less importance. In terms of techniques and procedures to be employed, this implies that the main consideration to apply when selecting or

[1] Although the topic is controversial, interested readers can refer to (Lodge 1981) – "if a novel did not bear some resemblance to other novel we should not know how to read it" – or (Anderson 2007) – "novelty (...) threatens to present a major obstacle to how and what these text can communicate to readers."– for more detailed discussion.

designing a narrative generation solution is to ensure that it will produce outputs that are clearly recognisable as instances of stories. A would-be designer of narrative generators needs to invest a significant portion of his time in developing techniques that will capture the unwritten laws of how to write, master the known tropes, and are aware of classic references in the field. This does not mean that novelty can be ignored. When reusing known structures new ways of instantiating them must be found or new twists added to give them a spark.

This need to strike a delicate balance between reproduction of known elements and innovation has traditionally been resolved by the application of knowledge-based approaches where knowledge is defined in meaningful units that capture the ingredients to be reproduced, and procedures are devised for recombining them into new material in a way that captures the essence of the craft without reproducing existing material literally. Finding the right balance is a significant challenge. If the reference material is represented with small granularity, local coherence of the results is assured but there may be insufficient knowledge to drive an interesting overarching plot. If the granularity is too large, the structure of the results does not depart from that of the reference material enough to suggest novelty.

The present paper reviews existing approaches to story generation focusing on the granularity at which they represent the knowledge that drives plot construction, and discusses their relative merits. Based on evidence arising from this analysis, a new plot representation format is proposed that combines the advantages of low and high granularity solutions, and a construction algorithm designed to guarantee the acceptability of output stories is presented.

Related Work

The work presented in this paper is inspired by and departs from some prior theoretical accounts of plot and some computational approaches to the representation and construction of narrative.

Some Relevent Theoretical Accounts of Plot

The most popular representation schema for plot among the early story generators has been the concept of *character functions* as presented in the morphology of the Russian folk tale (Propp 1968). A character function for Propp is an abstraction over certain actions of the characters that are relevant to the overall plot. Examples of these actions relevant for the plot are: performing a villainous act, starting a fight, winning a fight, departing on a journey or returning from a journey, but also less active choice points in the story such as deciding to take action in view of a villainy, recognising a character that was in disguise, or rewarding someone. For Propp, these character functions, in their abstract versions, were shared across the set of folk tales he analysed. He considered them grouped into a set of spheres, each one of them associated with one of the dramatis personae: the hero, the villain, the victim... Furthermore, Propp postulated an overarching canonical sequence that described the relative order in which the character functions appeared in the plot.

Computational solutions based on Propp's account have tended to borrow some parts of the analysis while forgetting others. The part of Propp's account most often used are his character functions, which correspond to representing the plot with a small granularity.

At the opposite extreme of the spectrum, there have been a number of efforts to represent plot in terms of its overall structure as a whole. These efforts postulated a number of abstractions of the structure of a plot that act as the set of master plots available for building stories. The number of such master plots varies between the single plot structure known as *the hero's journey* (Campbell, Cousineau, and Brown 1990) and Plotto's 1,462 plots (Cook 2011) with intervening values at twenty (Tobias 2012) or seven (Booker 2004). This diversity in values relates to the level of abstraction at which the plots are described. This in itself presents a serious challenge for a researcher hoping to establish an appropriate representation scheme, but it also suggests that a solution articulated at a lower level of granularity, which described complex plots as combinations of smaller units might be better suited to capture the complexity in the data in a more efficient way.

Although many more accounts of plot from a theoretical point of view have been reviewed to inform the present research, it is beyond the scope of this paper to list them all. But there are two specific ones that suggest a middle way to the dilema discussed so far. Working at different level of granularity in representing his material, Polti describes thirty six dramatic situations (Polti and Ray 1916) that can be used by a playwright to structure his material. On close perusal these situations are not full fledged plots. The examples provided in the book very rarely intend the situation in question as a description of the complete plot, and in describing them the author often refers to particular acts of the plays in question as instances of these situations. Although no definition of a situation is given, they are at a certain point referred to as "actions possible to the theater". It seems they are intended as building blocks for a plot. Some of them correlate reasonably well with certain of Propp's functions (abduction, pursuit) but others seem to operate at a larger degree of granularity than Propp's character functions (crime pursued by vengeance). Another important insight can be obtained from Forster's analysis of plot (Forster 1927). For Forster, plot is distinct from a chronological sequence of events in that the events within a plot must be connected by some kind of causality that drives the sequence and gives it meaning. These two sources suggest that a representation at an intermediate level of granularity, coupled with a procedure for explicitly connecting them to one another with some kind of motivational link, could provide a solid basis for representing plot.

Some Computational Representations of Plot

Existing efforts at representing plot with generative purposes may also be analysed in terms of the granularity at which they consider their knowledge units. In each case, they also consider additional procedures for governing how the individual units are combined into output plots.

There are many story generators that rely on planning solutions (Riedl 2004), based on the assumption that a plan – which draws a connecting path between an initial situation and a goal – has a certain parallelism with a story. In these approaches, the basic unit for representing plot is a planning operator, which involves a representation of an action or event, with associated preconditions and postconditions. Local coherence is achieved by the precondition/postcondition links between actions in a plot, and the overall structure is controlled by the definition of an initial situation and a goal that the plot will conform to. The planning approach to generating plot provides very strong local coherence and weak control over the story arc over a complete plot, so it has been a preferred solution for interactive storytelling, where allowing the user to interfere with the story at will with no obvious loss of coherence is actually a virtue. Attempts have been made to enhance the control over the structure of the plot by the use of *vignettes*, which are defined as “plot fragments that are a priori known to be good” (Riedl and Sugandh 2008; Riedl and León 2008). This approach involves adding an intermediate level of granularity to the representation of plot, one where a set of actions are already packed into a larger fragment that is reused in the construction process.

The Mexica system (Pérez y Pérez 1999) also relies on a combination of *story actions* – atoms of action defined with preconditions and postconditions– with larger structures called *story contexts* that determine how such story actions can be combined. Story contexts are extracted from a set of prior stories that established the knowledge base of the system.

A different family of narrative generators relies on story grammars (Rumelhart 1975) to represent story structure (Lang 1999; Bringsjord and Ferrucci 2000). Rummelhart’s original proposal for story grammars included a syntactic component – which establishes the sequence of plot tokens in the surface form – coupled with a semantic component – which imposes restrictions on conceptual relations that should hold between tokens adjacent in the plot sequence for the story to be meaningful. This combination matches Forster’s concept of plot as requiring conceptual connections across story elements. Computational approaches based on Rummelhart’s story grammars have attempted to implement the semantic component in different ways, such as additional Prolog clauses for the semantics (Lang 1999) or satisfaction of a theory of betrayal (Lang 1999; Bringsjord and Ferrucci 2000).

Propp’s formalism for representing plot in terms of characters functions has been employed repeatedly in historic systems. The most faithful rendering of Propp’s approach was developed in (Gervás 2015), where plots were represented as sequences of plot atoms represented as instances of Propp’s character functions, and combined based on preconditions and particular heuristics to capture long term dependencies between them. These long term dependencies were represented in the system as a specific additional knowledge resource.

Many of these approaches to plot generation have been compared in terms of their relative ability to satisfy specific metrics related to quality of the resulting plot structures (Gervás 2017). The conclusions of that comparison were that different approaches – and the associated representations – focus on particular features that are necessary in a story but not altogether sufficient. The combined set of features so identified was not covered by any of the individual approaches. The solution proposed in this paper attempts to capture a broad combination of the features in a single representation.

More recent work on the generation of fictional stories based on observed facts (Gervás 2018a) – referred to as *storification* – rely on a representation of plot as a sequence of *plot elements*, where each plot element is similar to a character function but explicitly holds additional information to indicate how the roles specific to the plot element (kidnapper, kidnapped) are filled in by roles that are relevant to the plot (villain, victim). In this work, plots are applied only as fixed set of schemas represented in this form, which are matched to observed facts to reach a joint representation that combines fact and fiction into a new story. An extension of this work (Gervás 2018b) first proposed the concept of *axis of interest* as a representational mechanism for articulating the representation of plots. In that approach, axes of interest were only used as representational devices, to allow for the construction – by hand – of a broader set of plot schemas to use as resources in the storification process.

Generating Acceptable Narrative Plots

Given that generation of narrative plots is such a complex task, one possible approach to simplify the engineering of a system might be to first model the ability to create acceptable narrative plots – in the same way as would-be authors first have to master the craft of generating acceptable stories – and later refine the approach to focus on “good” or “novel” stories. By concentrating on the simpler problem of generating acceptable stories (regardless of their novelty) the development task can focus on solving the difficulties involved in achieving acceptability. If a procedure is found to generate a broad range of acceptable stories, the achievement of quality and novelty may be attempted later by the application of metrics to filter the outputs. This would correspond to the *filtration* approach to the construction of generative systems, defined in (Ventura 2016) among the ones with an option for being considered beyond “mere generation”. The present paper focuses on the first stage of such an approach: the generation of a search space of acceptable narrative plots.

The approach followed to achieve this involves composing plot schemas as the interweaving of a number of linear substructures, themselves built up of conceptually interconnected plot atoms.

Representing Plot

The plots considered for the present paper are represented in terms of structured compositions of basic units called *plot atoms*. A plot atom is built along the lines of the plot elements defined in prior approaches to storification (Gervás 2018a; 2018b): a unit similar to a character function which explicitly holds additional information to indicate how the

AXISofINTEREST =	DONOR
AXISofINTEREST PROTAGONIST =	tested
AXISofINTEREST ROLES =	tested tester user gift
PLOT-SPAN-NAME =	Tested
Tested	characters(tested=X,tester=Y)
Character'sReaction	characters(tested=X,tester=Y)
ProvisionOfAMagicalAgent	characters(tested=X,tester=Y) objects(gift=Z)
PLOT-SPAN-NAME = UseOfAMagicalAgent	
UseOfAMagicalAgent	characters(user=X) objects(gift=Z)

Table 1: The Axis of Interest for the DONOR sequence (Propp 1968). Upper case letters indicate free variables.

roles specific to the plot element (kidnapper, kidnapped) are filled in by roles that are relevant to the plot (villain, victim). This refinement allows for interesting articulation between roles specific to a plot atom and more general roles that refer to the overall plot.

The plot atoms in a plot are organised in a complex structure that combines different sequences of plot atoms. This is required to allow for the concepts of plots that relate actions that take place at non-contiguous points in time (villainy early in the story, revenge at the end of it, with other unrelated events happening in between), or plots that combine more than one subplot (each subplot is a different sequence of plot atoms which may be interleaved with the other subplots by breaking their sequence down into smaller subsequences that constitute different scenes of the subplot). This type of complex structure is represented by a recursive data structure: the plot span. A *plot span* represents a span of plot, constituted by a sequence of plot atoms (or smaller spans). The idea is to capture the concept of a number of plot atoms appearing as a structural unit in a plot, but not necessarily occurring contiguously in the discourse for the plot. Plot spans of this type can be used to represent complete plots. When a plot span represents a complete plot, it is called a *plot schema*. Plot spans can also be used to describe intermediate units of plot structure that involve a set of plot atoms related by a long range dependency. These are called *axes of interest*. 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 atoms are structurally connected. An axis of interest has a set of narrative roles – those of its constituent plot atoms – that are initially free variables but which can be instantiated to specific constants when the axes of interest is combined into larger structures.

Two examples of axes of interest are shown in Tables 1 and 2. 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 can be combined together, weaving their corresponding subspans with those of other axes of interest, to form plot schemas. A *plot schema* encodes the way in

AXISofINTEREST =	CONFLICT
PROTAGONIST =	attacker
ROLES =	attacker defender winner looser
PLOT-SPAN-NAME =	Struggle
Struggle	characters(attacker=X,defender=Y)
PLOT-SPAN-NAME =	Victory
Victory	characters(winner=X,looser=Y)

Table 2: The Axis of Interest for CONFLICT

PLOT-SCHEMA =	OCM-DonFight	
PROTAGONIST =	hero	
DONOR	Tested	characters(tested=**hero**,tester=donor) objects(gift=gift)
CONFLICT	Struggle	characters(attacker=**hero**,defender=villain)
DONOR	UseOfAMagicalAgent	characters(user=**hero**) objects(gift=gift)
CONFLICT	Victory	characters(winner=**hero**,looser=villain)

Table 3: Example of plot schema for a basic plot combining axes of interest for DONOR and CONFLICT. The first column shows the interweaving of the axes of interest. Horizontal lines show the boundaries between spans corresponding to different axes of interest. The co-occurrence of constants on both sides of a boundary line in the final column – shown in **bold** – indicates the presence of a plot link at that point between the two axes.

which several axes of interest combine together to form the plot span for an elaborate plot. In a plot schema, the plot atoms from the axes of interest that have been combined appear in an ordered sequence that corresponds to the discourse for the plot schema, but each plot atom is labelled to indicate which axes of interest it corresponds to. Additionally, the plot schema lists for each plot atom how the roles specific to the various axes of interest are instantiated in terms of the set of constants that encode the overall set of narrative roles involved in the plot schema.

An example of plot schema is presented in Table 3, which shows how the DONOR and CONFLICT axes of interest – both abstracted from Propp's account of the Russian folk tale – are interleaved to form a very basic plot where the hero defeats the villain using a magical agent acquired earlier in the story. It also shows how the narrative roles for the plot (*hero*, *villain*, *victim*, *donor*) are mapped to the roles specific to the plot atoms of the constituent axes of interest (*tested*, *tester*, *user* for the DONOR axis of interest and *attacker*, *winner* for the CONFLICT axis of interest). This ensures that the various plot atoms in the plot are instantiated in a manner coherent with the narrative roles that the characters play in the overall plot schema.

Restrictions on Axes of Interest Combination

For a plot schema to be considered a valid narrative plot, the variables in the axes of interest that compose it must be instantiated in a coherent manner. For the example in Table 3, the required restrictions can be expressed by requiring that the character who acts as *tested* and *user* in the DONOR axis be the same one who acts as *attacker* and *winner* in the

DONOR CONFLICT	ProvisionOfAMagicalAgent Struggle	tested attacker
CONFLICT DONOR	Struggle UseOfAMagicalAgent	attacker user
DONOR CONFLICT	UseOfAMagicalAgent Victory	user winner

Table 4: Plot links for the connections between the DONOR and the CONFLICT axes of interest in the context of the plot schema in Table 3. Each plot link is described by two lines in the table. The first column indicates the axes of interest involved, the second column indicates the plot atoms appearing at the boundaries, and the third column indicates the roles in each axis that need to be instantiated to a a shared character. A plot link exists by virtue of the pairs of axis-specific roles in the third column being both instantiated with the same character at different sides of a boundary between two adjacent spans.

CONFLICT axis. Note that unless this condition holds the proposed plot schema makes no sense. These restrictions across axes of interest are captured in terms of *plot links* such as the one shown in Table 4.

If the plot schema is created by hand, the plot links can be abstracted from it by identifying shared variables co-ocurring at boundaries of adjacent spans. Plot links extracted in this fashion can then be used to guide recombination of existing axes of interests into new plot schemas according to the construction procedure described below.

A sequence of plot atoms from several axes of interest, interwoven into a single linear sequence, is considered *valid from a narrative point of view* if at any point in the sequence where plot atoms from different axes of interest appear contiguously, there exists a plot link connecting a shared variable.

A plot schema is considered valid if the sequence of plot atoms it encodes is valid.

The simplicity of representing plots in this way 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.

The search space of possible plots obtainable with this representation can be generated by exploring all combinations of the available axes of interest. A combination of two sequences of plot atoms is built by considering all possible interleavings of the plot atoms in them, and pruning any branches of this search where plot atoms from different axes of interest appear contiguously and are not supported by plot links.

Knowledge Engineering Issues

A set of plot atoms and axes of interest built by combining them needs to be crafted by hand. This is a significant knowledge engineering effort, along the lines of others previously described in the literature on narrative generation (Gervás, León, and Méndez 2015). For the sake of comparability, the set of basic plot atoms proposed in (Gervás, León, and Méndez 2015) was considered, together with a set of 19 axes of interest corresponding to the subsequences of Propp's canonical sequence (Propp 1968), and the elementary encoding of instances of Booker's seven basic plots (Booker 2004) proposed in (Gervás, León, and Méndez 2015).

The proposed representation has the advantage that the set of plot links can be mined from a set of instances of plot schemas built using this vocabulary of basic elements. For the purposes of this paper, an initial set of 34 plot schemas was constructed to act as seed. A set of plot links was created by means of a bootstrapping procedure:

- the set of seed plots was parsed to obtain an initial set of plot links
- the set of plot links so obtained was used to drive a construction procedure that generated a set of new plot schemas by exploring pairwise combinations of the existing axes of interest
- the set of output plot schemas was manually revised for correctness and expanded by the construction of further combinations analogous to the ones in the output
- the revised and expanded set of plot schemas was parsed to obtain further plot links

This procedure resulted in a set of 221 plot schemas, which gave rise to a set of 423 plot links between the 34 plot atoms over the 19 axes of interest.

Number of axes of interest combined	Number of plots generated
2	144
3	1,051
4	10,301

Table 5: Number of plots generated for combinations of different numbers of axes of interest.

Testing Generative Capacity

The resulting set of resources can then be used to generate combinations of larger numbers of axes of interest. Overall numbers of plots generated for different values of the number of axes of interest considered are given in Table 5.

The amount of outputs generated makes it impractical to carry out an exhaustive quality check. Random sampling was carried out by generating a number at random within the range of numbered outputs and checking the corresponding plot. The examples presented below have been chosen in this fashion. Overall they seem to be acceptable as possible instances of plots, and some of them are actually reasonable in the sense that one can follow a certain logic in the way that the plot atoms follow one another.

The system generates conceptual descriptions of plot schemas much along the lines of the examples of knowledge resources presented earlier. For ease of reading, a simple template-based text realizer has been developed that converts such conceptual representations into readable text. The

RIVALRY	Rivalry	**Hero** develops rivalry with shadow.
CROSSDRESSING	CrossDressing	**Hero** dresses up as a member of the opposite sex.
PURSUIT	Pursuit	**Hero** is pursued by villain.
RIVALRY	Cooperation	**Hero** cooperates with shadow.
RIVALRY	RivalReconciliation	**Hero** ends rivalry with shadow.
PURSUIT	RescueFromPursuit	**Hero** avoids pursuit.
CROSSDRESSING	Recognition	**Hero** is recognised.

Table 6: Example output for combination of 3 axes of interest: RIVALRY, PURSUIT and CROSSDRESSING.

RAGS2RICHES	Poverty	Hero suffers poverty.
RAGS2RICHES	Aspiration	**Hero** has aspiration.
RELENTINGGUARDIAN	CoupleWantsToMarry	**Hero** wants to marry love-interest.
RELENTINGGUARDIAN	UnrelentingGuardian	**Obstacle** objects to proposed union of **hero** with love-interest.
TASK	DifficultTask	**Hero** is set a difficult task by **obstacle**.
TASK	Solution	**Hero** solves the task.
RAGS2RICHES	Transformation	**Hero** is transformed.
RELENTINGGUARDIAN	RelentingGuardian	**Hero** convinces obstacle
RELENTINGGUARDIAN	Wedding	**Hero** marries love-interest.
RAGS2RICHES	Reward	**Hero** is rewarded.

Table 7: Example output for combination of 3 axes of interest: RAGS2RICHES, TASK and RELENTINGGUARDIAN.

quality and elaboration of the resulting texts has deliberately been kept low, to avoid confusion between any merits arising from the conceptual structure of the narrative plots generated and any beauty that may arise from the texts generated to render these structures.

Results for different combinations of 3 axes of interest are shown in Tables 6, 7 and 8. In each case the first two columns show the axes of interest and the plot atoms involved respectively, and the third columns shows the output story rendered as text. Boundaries between plot spans corresponding to different axes of interest are indicated by horizontal lines. The characters instantiating the plot links across the boundaries are shown in **bold**.

Tables 9 and 10 show examples for different combinations of 4 axes of interest.

Because they have been obtained via random sampling, the examples presented here should not be considered as selected outputs, but rather as samples out of a very large search space that can be generated by the proposed solution. Because of the high number of potential plot links available, the current procedure generates more than one possible interweaving for a given choice of axes of interest, depending on the choices made over the available plot links. Not all of them are equally fortuitous in terms of the story they represent. For instance, the example in Table 9 might have been a better story if the recognition of the validation received had occurred not immediately after it had been granted, but at a later point in the story, maybe after the hero had cooperated with the shadow. Such a combination may have occurred in the set of outputs, but checking them by hands to find it is a serious task.

SHIFTINGLOVE	BoyMeetsGirl	Hero meets and starts relationship with love-interest.
SHIFTINGLOVE	LoveShift	**Hero** looses the atention of love-interest.
CONFLICT	Struggle	**Hero** fights with villain.
CONFLICT	Victory	**Hero** achieves victory over villain.
REPENTANCE	Transformation	**Hero** is transformed.
REPENTANCE	Repentance	**Hero** repents.
SHIFTINGLOVE	Reconciliation	**Hero** makes up with love-interest.
REPENTANCE	RepentanceRewarded	**Hero** sees repentance rewarded.

Table 8: Example output for combination of 3 axes of interest: SHIFTINGLOVE, CONFLICT and REPENTANCE.

RAGS2RICHES	Poverty	Hero suffers poverty.
RAGS2RICHES	Aspiration	**Hero** has aspiration.
RIVALRY	Rivalry	**Hero** develops rivalry with shadow.
CROSSDRESSING	CrossDressing	**Hero** dresses up as a member of the opposite sex.
VALIDATOR	Tested	**Hero** is tested by validator.
VALIDATOR	Character'sReaction	Hero reacts to the test by validator.
VALIDATOR	Validation	Hero is validated by validator.
VALIDATOR	ValidationRecognised	**Hero** sees validation recognised.
RAGS2RICHES	Transformation	**Hero** is transformed.
RIVALRY	Cooperation	**Hero** cooperates with shadow.
CROSSDRESSING	Recognition	**Hero** is recognised.
RIVALRY	RivalReconciliation	**Hero** ends rivalry with shadow.
RAGS2RICHES	Reward	**Hero** is rewarded.

Table 9: Example output for combination of 4 axes of interest: RAGS2RICHES, RIVALRY, CROSSDRESSING and VALIDATOR.

SHIFTINGLOVE	BoyMeetsGirl	Hero meets and starts relationship with love-interest.
SHIFTINGLOVE	LoveShift	Hero looses the atention of **love-interest**.
JOURNEY	Departure	**Love-interest** departs.
TASK	DifficultTask	Dispatcher sets a difficult task to **love-interest**.
PURSUIT	Pursuit	**Love-interest** is pursued by villain.
PURSUIT	RescueFromPursuit	**Love-interest** avoids pursuit.
TASK	Solution	**Love-interest** solves the task.
JOURNEY	Return	**Love-interest** returns.
SHIFTINGLOVE	Reconciliation	Hero makes up with **love-interest**.

Table 10: Example output for combination of 4 axes of interest: TASK, PURSUIT, SHIFTINGLOVE and JOURNEY.

Overall, the sampling carried out over the output suggests that there are no seriously flawed results. Some of the stories are more interesting than others, but this is to be expected in an exhaustive enumeration of the search space afforded by the chosen representation. As it stands, the proposed procedure constitutes a valuable initial module over which processes of filtering or selection may be developed by the introduction of appropriate metrics designed to capture particular concepts of story quality.

Discussion

The proposed system is discussed in terms of some of its shortcomings and in terms of its relation to prior work.

Aspects in Need of Improvement

The experiments carried out in support of the writing of this paper have uncovered a number of issues that may need atention.

The current approach to the representation of connections between plot atoms within a plot schema is based on the identification of shared variables representing characters that participate in each of the plot atoms with a different role. This is adequate for capturing a significant number of connections, but some connections very relevant to traditional stories are beyond the reach of the formalism. For instance, many classic plots are initiated by a villainy that results in the hero being faced with the call to action (Propp 1968; Campbell, Cousineau, and Brown 1990). The current version of the formalism cannot capture this kind of connection, because the hero is not explicitly mentioned in the villainy and the villain is not explicitly present in the call to action.

So no plot link can be established between them. This occurs in a number of further instances of plot where there is a conceptual connection between adjoining plot atoms that is relevant to the structure of the story but does not rely on co-occurrence of any character across the pair. This is one of the reasons why the number of generated plots falls below the expected number of possible combinations of the resources invested – the set of plot links mined from 221 seed plots produces only 144 constructed plots, as described above. The underlying formalism for representing plots may be revised in the future to address this limitation.

Another important shortcoming of the proposed formalism in its current version is that, due to the limitations of the initial implementation of the procedure for instantiating variables when weaving plot spans together, it does not allow for an axis of interest to be used more than once in a given plot. As a result, the system cannot represent a number of classic stories, such as action tales involving more than one fight, fairy tales where the hero attempts a task three times before succeeding, or comedy plots with multiple romantic couples switching affinities between them. This is a significant handicap that needs to be overcome in the future.

Relation to Prior Work

The proposed representation for plot shows significant parallelism with some of the approaches reviewed in the section on *Computational Representations of Plot*.

The decision to enrich the representation of plot atoms with explicit indication of the characters that take part in them, expressed in terms of the roles these characters play in the overall narrative structure of the plot, operationalises the concept of *sphere of action of the dramatis personae* as described by Propp.

It is interesting to see that Propp's overall schema for the representation for plot already includes three levels of representation of plot that match closely those that are proposed in the present paper: one of atoms to be recombined, one of relations between the atoms in terms of characters that play the fundamental roles in them, and one of relative ordering within the general plot arc.

The need for the additional mechanism of plot links to capture a conceptual structure across the plot atoms that is different from and goes beyond their adjacency relations within the sequence of the plot is inspired by Forster's insight that plot incorporates additional levels of connection beyond simple chronological sequence (Forster 1927).

The abstraction of a unit for the representation of plot that is intermediate between a full plot and the kind of plot atom illustrated by Propp's character functions – or plan operators understood as story actions – was already present in Polti's *dramatic situations* (Polti and Ray 1916). Because they sometimes refer to ingredients of the plot structure that span across its length – such as *crime pursued by vengeance* – they have close similarity with the proposed concept of axis of interest.

The use of the concept of an axis of interest to tie together a number of plot atoms into a construction unit larger than an atom but smaller than a full plot has similarity to that of *vignettes* as proposed in (Riedl and Sugandh 2008; Riedl and León 2008). However, vignettes tend to correspond to short sequences of consecutive actions that fill a single gap at a particular point of time in the plot, whereas an axis of interest will usually span two different moments in time that are relevant for a plot, in a way that allows the encoding of long range dependencies between separate moments of the plot.

By their construction and the way in which they are combined to construct plot schemas, axes of interest are designed to capture conceptual dependencies between plot atoms that are conceptually connected – like an imprisonment and the release of the prisoner – but occur at places in the story distant from one another. This is a significant advantage in that it allows for the construction of structurally complex stories spanned by conceptual links between distant elements.

This use of axes of interest and plot links as extra layers of meaning over the order in which the plot atoms occur in the discourse is related to the need identified in (Rumelhart 1975) to represent the structure of stories at both a more superficial syntactic level and a deeper semantic level. Some of the shortcomings of the solution proposed in this paper arise from the fact that the current representation of conceptual connections in terms of plot links is still too close to the syntactic level represented by explicit mention of a character in a plot atom.

The proposed bootstrapping solution for exploiting the generator itself to construct plot schemas that are then adapted to provide further sources for mining knowledge for the system follows an existing line of work on reducing the bottleneck of knowledge acquisition. Prior solutions to the task of engineering the knowledge resources for story generators had been proposed based on mining crowdsourced plot graphs (Li et al. 2013) or applying Qualitative Knowledge Engineering methodologies (O'Neill 2013). The procedure for extracting the knowledge resource for driving connections between plot atoms – in the case of this paper, plot links – by parsing instances of stories as represented within the system – here, plot schemas – is borrowed from the way the Mexica system parses prior stories to build its story contexts.

Conclusions

The system described in this paper presents significant advantages in terms of how it captures the conceptual complexity of plot, how it can be used to construct a useful number of knowledge resources required for operation, and how it allows the construction of relatively large number of acceptable plots.

In its current version the system is not intended as a fully creative plot generation system, but rather as a prototype for the initial stage of a developing system. This initial stage would address the task of generating a broad range of acceptable stories instead of aiming for a small set of stories of high quality. In a way, it is intended to model the craft of putting together something that can be recognised as a story, not necessarily a good one.

Even within this scope, the work reported in this paper indicates that some engineering challenges remain unresolved. A number of shortcomings have been identified and further work on the system will hope to address them to improve both the range and the acceptability of system outputs.

The vocabulary of basic resources – both in terms of plot atoms and in terms of axes of interest – may be extended to increase the expressive power of the proposed representation.

Once those planned improvements have been carried out, long term future work may consider the development of additional modules designed to identify parameters that relate to the perceived quality of stories. If such modules become available, more elaborate procedures may be designed that start to consider the generation a smaller number of stories of better quality, or that focus on generating stories that are significantly different from the stories already known to the system. Consideration of the creativity of the system as a story generator would need to wait upon the development of these additional modules for quality and novelty metrics, and the design of this specialised story generation system.

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

Anderson, E. H. 2007. Novelty in Novels: a Look at What's New in Aphra Behn's 'Oroonoko'. *Studies in the Novel* 39(1):1–16.

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

Bringsjord, S., and Ferrucci, D. A. 2000. Artificial intelligence and literary creativity: Inside the mind of brutus, a storytelling machine. *Computational Linguistics* 26(4).

Campbell, J.; Cousineau, P.; and Brown, S. 1990. *The Hero's Journey: Joseph Campbell on His Life and Work*. Collected works of Joseph Campbell. New World Library.

Cook, W. 2011. *Plotto: The Master Book of All Plots*. Tin House Books.

Forster, E. M. 1927. *Aspects of the novel*. New York: Harcourt.

Gervás, P.; León, C.; and Méndez, G. 2015. Schemas for Narrative Generation Mined from Existing Descriptions of Plot. In Finlayson, M. A.; Miller, B.; Lieto, A.; and Ronfard, R., eds., *6th Workshop on Computational Models of Narrative (CMN 2015)*, volume 45 of *OpenAccess Series in Informatics (OASIcs)*, 54–71. Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

Gervás, P. 2015. Computational Drafting of Plot Structures for Russian Folk Tales. *Cognitive Computation*.

Gervás, P. 2017. Comparative evaluation of elementary plot generation procedures. In *6th International Workshop on Computational Creativity, Concept Invention, and General Intelligence*.

Gervás, P. 2018a. Storifying observed events: Could i dress this up as a story? In *5th AISB Symposium on Computational Creativity*. University of Liverpool, UK: AISB.

Gervás, P. 2018b. Targeted storyfying: Creating stories about particular events. In *Ninth International Conference on Computational Creativity, ICCC 2018*. Salamanca, Spain: Association of Computational Creativity.

Lang, R. R. 1999. A declarative model for simple narratives. In *Proceedings of the AAAI Fall Symposium on Narrative Intelligence*, 134–141. AAAI Press.

Li, B.; Lee-Urban, S.; Johnston, G.; and Riedl, M. O. 2013. Story generation with crowdsourced plot graphs. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence*, AAAI '13.

Lodge, D. 1981. *Working with Structuralism: Essays and Reviews on Nineteenth and Twentieth Century Literature*. Routledge & K. Paul.

O'Neill, B. 2013. *A Computational Model of Suspense for the Augmentation of Intelligent Story Generation*. Ph.D. Dissertation, Georgia Institute of Technology, Atlanta, Georgia.

Pérez y Pérez, R. 1999. *MEXICA: A Computer Model of Creativity in Writing*. Ph.D. Dissertation, The University of Sussex.

Polti, G., and Ray, L. 1916. *The Thirty-six Dramatic Situations*. Editor Company.

Propp, V. 1968. *Morphology of the Folktale*. Austin: University of Texas Press.

Riedl, M., and León, C. 2008. Toward vignette-based story generation for drama management systems. In *Workshop on Integrating Technologies for Interactive Stories - 2nd International Conference on Intelligent Technologies for Interactive Entertainment*.

Riedl, M. O., and Sugandh, N. 2008. Story planning with vignettes: Toward overcoming the content production bottleneck. In *Interactive Storytelling*, volume 5334 of *Lecture Notes in Computer Science*, 168–179. Springer.

Riedl, M. 2004. *Narrative Planning: Balancing Plot and Character*. Ph.D. Dissertation, Department of Computer Science, North Carolina State University.

Ritchie, G. 2007. Some empirical criteria for attributing creativity to a computer program. *Minds and Machines* 17:76–99.

Rumelhart, D. E. 1975. Notes on a schema for stories. *Representation and Understanding: Studies in Cognitive Science* 211–236.

Tedford Jones, J. 2002. Depending on memory: Intertextuality in popular fiction. *The Journal of American Culture* 25:81–84.

Tobias, R. 2012. *20 Master Plots: And How to Build Them*. F+W Media.

Ventura, D. 2016. Mere generation: Essential barometer or dated concept? In *Proceedings of the Seventh International Conference on Computational Creativity (ICCC 2016)*. Paris, France: Sony CSL.