Pruning Worlds into Stories: Affective Interactions as Fitness Function*

Pablo Gervás[0000-0003-4906-9837] and Gonzalo Méndez[0000-0001-7659-1482]

Facultad de Informática, Universidad Complutense de Madrid, Madrid, 28040 Spain
{pgervas,gmendez}@ucm.es
http://nil.fdi.ucm.es

Abstract. An important challenge when trying to find a story to tell about some set of events that has already happened is to identify the elements in that set of events that will make a story that moves the intended audience. One possible criterion is to consider events that involve significant changes in the emotional relations between the characters involved. The present paper explores a computational model of this particular approach to the task of storytelling. An evolutionary solution is used to explore the logs of an agent-based social simulation, using metrics on the evolution of affinity between characters as fitness function, to identify sequences of events that might be good candidates for moving stories.

Keywords: story generation · affective interactions · evolutionary approach.

1 Introduction

The ability to tell stories originates as means for telling a target audience about something that has happened. But an important restriction operates: the telling of the story must require significantly less time than the original events took to happen. So building the story involves an effort to select from the set of events observed the optimal subset that conveys the "meat" of the story. In most cases, additional restrictions are applied to ensure that the telling of the story achieves specific purposes. A very common purpose is to ensure that the story conveyed have some emotional impact on the audience.

The present paper explores a computational model of the task of building a story from an observed set of events, with the particular purpose to maximise the emotional impact on the audience. The set of observed events to consider is taken to be the log of an agent-based social simulation in which agents engage in interactions intended to lead to romantic entanglements between them.

* This paper has been partially funded by the projects CANTOR: Automated Composition of Personal Narratives as an aid for Occupational Therapy based on Reminiscence, Grant. No. PID2019-108927RB-I00 (Spanish Ministry of Science and Innovation) and the ADARVE (Análisis de Datos de Realidad Virtual para Emergencias Radiológicas) Project funded by the Spanish Consejo de Seguridad Nuclear (CSN).

A set of metrics is designed to compute quantitative measures related to the affective relationships between the agents. These metrics are employed to drive the fitness function for an evolutionary search that identifies subsets of the original set of events that maximise intensity, contrast and evolution in the affective relationships between the agents featured.

2 Related Work

Three topics are considered relevant for this paper: hypotheses linking stories with the evolution of affect between characters, computational approaches to constructing stories from observed sets of events, and evolutionary solutions for building stories.

2.1 Stories and Emotion

The importance of the affective meaning of a story–in contrast to other genres such as essays–is emphasised by Egan [3]. In his analysis of plot, Egan states that in constructing a story "events have, nevertheless, been selected and juxtaposed in order to achieve a particular emotional response and assert a particular sense of causality".

An example of a computational construction of stories that exemplifies this approach is the Mexica system [17], in which emotional information about relations between characters is used as driving force in construction of stories. Mexica operates over a representation that associates to each event details on both preconditions and effects in terms of affinities between characters. Its construction algorithm is designed to build stories that show significant evolution of the affinities between characters through the course of the story.

The idea that the evolution of particular values through the course of the story is a valuable feature in stories has also been defended by Weiland [23]. Weiland attributes particular importance to the evolution of particular characters, in what the defines as *character arcs*. These represent the evolution of the character through the story.

A similar concept is considered by [18], who use machine learning techniques over a corpus of 1,327 stories from Project Gutenberg's fiction collection to identify a set of six core *emotional arcs* that recur over the plots in the corpus. They postulate that these emotional arcs constitute "essential building blocks of complex emotional trajectories." The six emotional arcs in question are:

- *Rags to riches* (affect rises from start to end)
- *Tragedy*, or *Riches to rags* (affect falls from start to end)
- *Man in a hole* (affect falls then rises)
- *Icarus* (affect rises then falls)
- *Cinderella* (affect rises then falls then rises again)
- *Oedipus* (affect falls then rises then falls again)

They further study a corpus of movies in search of correlations between presence of these emotional arcs and interest demonstrated on the movie–in terms of downloads over time– and they conclude that Icarus (affect rises then falls), Oedipus (affect falls then rises then falls again), and two sequential Man in a hole arcs (affect falls then rises then falls then rises again), are the three most successful emotional arcs.

A more recent study explores the construction of distributed representations of character networks in stories using neural networks to create representations in terms of fixed-length vectors [12]. This solution provides a very powerful means of representing character networks that allows comparability across representations of stories of different lengths.

2.2 From Worlds to Stories: Computational Approaches

There are two different research lines focused on obtaining stories from representations of worlds: approaches based on selecting events from worlds to obtain stories of particular types, and approaches based on building narrative discourses to describe subsets of events from a given world.

The significant growth of realistic simulations of worlds arising from the rise of videogames led to efforts aimed at sifting through the material generated by such simulations in search for stories. This approach has been described as *curating storyworlds* and as *story sifting* [20]. Story sifting involves processing significant volumes of events produced by a simulated storyworld to identify subsets that match particular specifications of types of stories. There is often some type of request by a user that establishes which type of story is desired. Initial approaches to the task required that the user specify their request in some kind of technical language [1, 10, 20, 9]. Specific languages have been proposed to describe the *story sifting patterns* that may be used in these searches [9]. More recent approaches take advantage of neural representations and allow the user to establish a curve that describes the story arc they desire, and use complex algorithms to compare this curve to those obtained from potential selections from the search space [13]. When storyworld are large,[1] sifting in search of specific patterns may still yield large volumes of stories. To address this problem, statistical criteria to identify stories that appear less frequently has been proposed as ad additional filter [11].

There is a research line that focuses on the task from the specific point of view of how to convert a set of events from a storyworld–which may have taken place over a range of locations and time periods–into a linear sequence of facts, known as *narrative discourse* in such a way that: (1) all the relevant events are told, (2) the relative order in which they appear in the discourse allows the reader to perceives easily any relations between them, and (3) particular effects on the reader may be accomplished. The task is described as *composition of narrative discourse* [4]. Efforts in this research line also operate over an input that is some

[1] Searches through outputs of the Bad News game required processing 140 years of simulation of a unique procedurally generated small American town [21].

kind of log of a complete set of events, and select from those events a subset to include in the discourse. However, in this case, the selection is very specifically tied to the construction of the discourse, and it may involve complex decisions to modify either the granularity at which events are described or the relative order in which they are told. One particular interesting example informs the selection of events to include in the final discourse by potential matches between the resulting discourse and a set of possible plot schemas [5]. In this case, the plot schemas operate in a similar way to the story sifting patterns described above.

2.3 Evolutionary Solutions for Building Stories

Evolutionary solutions for building stories often combine a genetic representation that describes a combination of basic story-building units, a set of evolutionary operators over that genetic representation, and a fitness function constructed in terms of metrics of the individuals in the population interpreted as stories. Existing approaches differ in terms of what they consider basic story-building units and in terms of how they define their fitness functions. Instances of basic story-building units considered are: partially ordered graphs of events associated with particular entities [15], instantiations of plot-relevant units of abstraction [6], plot templates [8], or even full story drafts built by an automatic story generator [22]. Criteria used to develop fitness functions include: metrics on story coherence and story interest [15], causal relations between events and certain characters being involved in related events [6], metrics on semantic consistency of the story [8], or knowledge-based heuristics [22].

Of particular interest is the solution for generating small quests for games presented in [14] which evolves a set of plans–each representing a possible quest–using as fitness function the degree of matching between the tensions in the story and a target curve of evolving tensions provided as input.

More recent approaches have considered the importance of avoiding problems of non-synonymous redundancy and low locality [19] when devising genetic representations for stories [7].

3 Finding the Most Emotional Chains of Events in Relationships

Four aspects need to be considered: the agent-based simulation used as source, a genetic representation for stories in this context, metrics for measuring the affective impact of stories, and the evolutionary process for searching for affective stories.

3.1 The Charade Agent-based Simulator

Charade [16] is a multi-agent simulation designed to express relations and interactions between characters in a storyworld based on the existing affinities between the characters, and to model the evolution of these affinities through a given period of time. The characters may have different kinds of relationships between them (mates, friends, foes or indifferent), which may change in accordance to the evolution of their affinity level, as a result of the interactions (or lack of them) that take place between characters. The result is a log of interactions and evolutions of affinity levels which are subsequently used to generate episodes within a narrative [2].

An example of a fragment of a log for the Charade sytem is shown in Table 1. The simulation is run with 15 agents who do not all know each other. Each agent may or may not have a partner, a small set of friends (between 2 and 4) and may or may not have any enemies (1 or 2 at the start). Currently, spatial relations between the agents are not considered at all, nor are there any elements that can be interacted with. Interactions are driven by affinities between characters, and also act upon them. Probability of interaction is highest for partners, lower for friends, and lowest for enemies. Acceptance of proposals raises affinity between the characters, rejections and inactivity lower it. Runs are stopped when the logs reaches 2,500 events.

```
Megan PROPOSE friend_have_lunch Meredith
Lester PROPOSE friend_chat Robert
Suzette PROPOSE friend_chat Silvy
Betty PROPOSE friend_weekend_out Clark
Meredith PROPOSE mate_watch_tv Lester
Clark REJECT-PROPOSAL friend_weekend_out Betty
Lester REJECT-PROPOSAL mate_watch_tv Meredith
Meredith ACCEPT-PROPOSAL friend_have_lunch Megan
Lester affinity with Meredith 87
Violet PROPOSE friend_chat Megan
Clark affinity with Betty 67
Robert REJECT-PROPOSAL friend_chat Lester
Meredith affinity with Megan 72
Silvy ACCEPT-PROPOSAL friend_chat Suzette
Robert affinity with Lester 72
Betty affinity with Clark 50
(...)
```

Table 1. A example of a fragment of the log generated by the Charade multi-agent simulation system.

The set of events in the log is read into a conceptual representation to allow further processing.

The simulation operates as a distributed system, with each agent responding to the general evolution of the simulation both in terms of actions performed–whether proactively proposing interactions to other agents or reacting to proposals received–and in terms of affective responses–changes in affinity towards other agents. The sequence of events as it appears in the log lists facts in the chronological order in which they have been added to the simulation. From a narrative point of view, the relations between facts corresponding to a particular interaction–proposal, reaction to the proposal, impact of the interaction on the affinities between the characters involved– constitute the basic units of story.

In the pre-processing stage, the log is parsed to group together the sets of facts relate to a particular interaction into a single unit of this type.

An example of the parse of a Charade log into a set of single narrative units is shown in Table 2.

```
PLOT-PROJECTION 0
ProposeActivity {activity=friend_weekend_out, proposee=Clark, proposer=Betty}
PLOT-PROJECTION 1
ProposedActivityAccepted {activity=friend_weekend_out, proposee=Clark, proposer=Betty}
AffinityChange {triggerer=Clark, perceiver=Betty, impact=76}
AffinityChange {triggerer=Betty, perceiver=Clark, impact=51-->54}
PLOT-PROJECTION 2
ProposeActivity {activity=mate_go_to_cinema, proposee=Mary, proposer=Clark}
PLOT-PROJECTION 3
ProposedActivityRejected {activity=mate_go_to_cinema, proposee=Mary, proposer=Clark}
AffinityChange {triggerer=Mary, perceiver=Clark, impact=95}
AffinityChange {triggerer=Clark, perceiver=Mary, impact=84}
(...)
```

Table 2. A example of a fragment of the parse of a Charade log into a set of plot projections.

Affinities between two agents A and B are directed, so what A feels for B may differ from what B feels for A. They are represented on a scale between 0 and 100, with 0 representing strong dislike and 100 representing passionate love. The Charade system considers a classification of relations between agents in terms of the affinities between them:

foe affinity between 0 and 40
neutral affinity between 40 and 60
friend affinity between 60 and 80
mate affinity between 80 and 100

The type of relation that holds between two agents determines the subset of activities that they may consider together.

The behaviour of agents is informed by the affinities between them, and the reactions of agents alter the affinities between them. Although in principle all actions by the agents might trigger a change in affinity, in the current set up only the reactions to a proposal produce changes in affinity.

3.2 Genetic Representation of Stories

A story in our approach is considered to be a selection of a subset of the single narrative units that have been constructed from the log. Each individual in our populations will be one such story. Because the events are ordered chronologically in the log of the simulation, a relative ordering can be established between all the single narrative units by assigning to each of them the chronological time of the events in the single narrative unit that constitutes a reaction to the corresponding proposal. The existence of this relative order makes it possible to represent each potential story as a numerical vector that encodes which of the single narrative units is to be included in the story. At each position, a value of 1 indicates that the corresponding single narrative unit is to be included in the story, and a value of 0 that it is not.

3.3 Metrics for Affective Interactions

To establish quantitative metrics on the affective interactions featured in a story, we first construct a set of affective threads. Each *affective thread* is a numerical vector of the same length as the story, and which is made up of the numerical values of the affinity that a character A holds at each point of the story for another character B. It is therefore a numerical vector, as long as the story, of values between 0 and 100.

For a given story, the sequence of affinities between each character and the others is compiled, so that metrics can be computed for each affective thread.

The following features are considered relevant for measuring the evolution of affect through the duration of the story, as featured in a specific affective thread:

- the maximum distance of affinities shown from the neutral value (50)
- whether the affinities change polarity–from positive to negative or negative to positive–at some point
- whether the affinities change into an interval that defines a different type of relation between the characters

The following metrics are defined in terms of these features, by establishing different types of constraints over the values of the metrics for the complete set of affective threads in a story:

- highest ratio between number of affinity changes and thread length
- average distance from neutral affinity
- percentage of threads that transition across meaningful relation types at least once
- percentage of threads that change polarity at least once

3.4 Evolutionary Optimisation of Subset of Told Events

The evolutionary process operates over a given log for the Charade system. An initial population of story drafts of size N is built by assigning the value 1 to N positions chosen at random of the genetic vector that represents a draft. The system is initialised with proportional sets of story drafts for a number of given lengths.

Mutation is carried out by deciding at random whether to activate or deactivate a gene, and then selecting at random a gene of the corresponding type to change its state.

Crossover is carried out by selecting at random a position in the genetic vector, dividing each of the participating vectors at that point, and interchanging the corresponding constituent parts.

Both mutation and crossover operators can change the length of the story drafts they operate on, so some means may be required to control the size of the drafts in the population in case the metrics affective interactions favour drafts either longer or shorter length.

The basic score assigned to the affective quality of a story–a number between 0 and 100–is established as the average value of the four metrics presented in Section 3.3. More complex combinations may be considered as further work.

A fitness function is defined that weighs the basic score with two additional scores designed to control the type of output obtained: (1) a score on story length that assigns 100 to drafts within a given range of length, and 0 otherwise, and (2) a score on character variety that assigns 100 to drafts with number of characters within a desired range of values, and 0 otherwise.

Selection of individuals to pass onto the next generation is done applying accumulated fitness to ensure broad enough coverage of the search space.

3.5 Template Based Rendering of Events

In order to make the outputs of the system easier to understand, the set of events in a particular drafts has been rendered into text using simple templates to express the corresponding events as simple text. This is not intended to be the final form of expression of system outputs when employed for any practical purpose.

Consideration of state of the art techniques [24] for rendering as fluent prose the conceptual representations of story that the system produces is considered beyond the scope of the present paper, as it would not involve evolutionary techniques.

4 Discussion

The results of the proposed system are presented and the relation of the proposed approach with previous work is discussed.

For basic testing of the evolutionary procedure, the parameter on character variety is set to a specific value of two characters in a story–to test the system's ability for identifying interesting developments in the relationship between two characters–and the upper bound on draft length is set to 20.

Table 3 shows an example of system output expressed in terms of the internal representation format, together with the corresponding story rendered as text automatically by the system using basic templates for each of the actions involved. The evolutionary solution was run for 30 generations, with a population size of 20 individuals. The final score for this draft is 61/100.

One can see over this example the main features of the system in operation. Figure 1 shows the evolution of affinities between the characters in the example

PA-0-ProposeActivity-95	mate_dinner_with_candles / Megan / Tony
PA-0-ProposedActivityAccepted-98	mate_dinner_with_candles / Megan / Tony
PR-0-ProposeActivity-106	mate_dinner_with_candles / Tony / Megan
PR-0-ProposedActivityRejected-112	mate_dinner_with_candles / Tony / Megan
PA-1-ProposeActivity-377	friend_serious_talk / Megan / Tony
PA-1-ProposedActivityAccepted-378	friend_serious_talk / Megan / Tony
PA-2-ProposeActivity-409	friend_help / Tony / Megan
PA-2-ProposedActivityAccepted-410	friend_help / Tony / Megan
PR-1-ProposeActivity-801	friend_have_coffe / Tony / Megan
PR-1-ProposedActivityRejected-802	friend_have_coffe / Tony / Megan
PA-3-ProposeActivity-872	friend_serious_talk / Megan / Tony
PA-3-ProposedActivityAccepted-873	friend_serious_talk / Megan / Tony
PA-4-ProposeActivity-912	friend_chat / Tony / Megan
PA-4-ProposedActivityAccepted-913	friend_chat / Tony / Megan

Tony proposes to Megan to dinner with candles as mates. Megan accepts Tony's invitation to dinner with candles as mates. Megan proposes to Tony to dinner with candles as mates. Tony rejects Megan's invitation to dinner with candles as mates. Tony proposes to Megan to serious talk as friends. Megan accepts Tony's invitation to serious talk as friends. Megan proposes to Tony to help as friends. Tony accepts Megan's invitation to help as friends. Megan proposes to Tony to have coffe as friends. Tony rejects Megan's invitation to have coffe as friends. Tony proposes to Megan to serious talk as friends. Megan accepts Tony's invitation to serious talk as friends. Megan proposes to Tony to chat as friends. Tony accepts Megan's invitation to chat as friends.

Perceiver	Target	
Tony	Megan	85 85 81 81 74 74 78 78 72 72 74 74 76
Megan	Tony	96 96 96 96 82 82 80 80 80 80 82 82 78

Overall	Polarity Change	Distance to Neutral	Relation Change	Rate of Change
61	0	40	100	42

Table 3. Example of story draft obtained by story sifting from a Charade log followed by the automated template-based rendering of the story. For ease of understanding the sequence of values for each of the characters is shown below, followed by the scores for the various metrics.

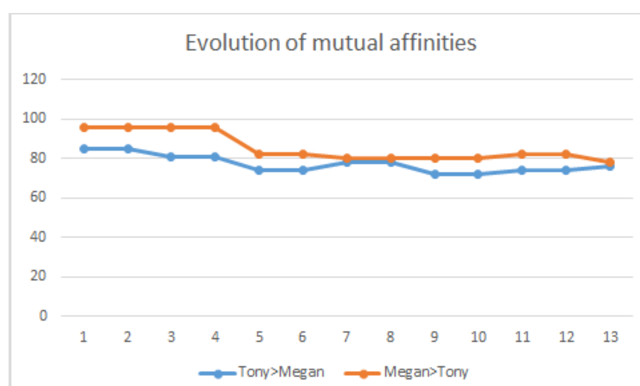

Fig. 1. Plot of the evolution of affinities between Megan and Tony. Vertical axis shows value of affinity, horizontal axis shows point in the discourse sequence for the draft.

in Table 3. An initial close relation progressively sours as Tony rejects Megan's proposals (at times 4 and 10), and in spite of a slight recovery when he accepts her offer of help (at time 8). The metrics show that there has been no polarity change (*Polarity change = 0*: in spite of the deterioration of their relationship, their affinities to one another remain positive throughout), that we are seeing a story that involves a certain passion (*Distance from neutral affinity=40*: at least at the beginning the affinities between them are high), that the relatioship between them is evolving (*Relation change=100*: Tony's affinity for Megan drops below the threshold of 80 at time point 5, making her more of a friend than a mate according to system rules), and affinities do change through the story (*Rate of change=40*).

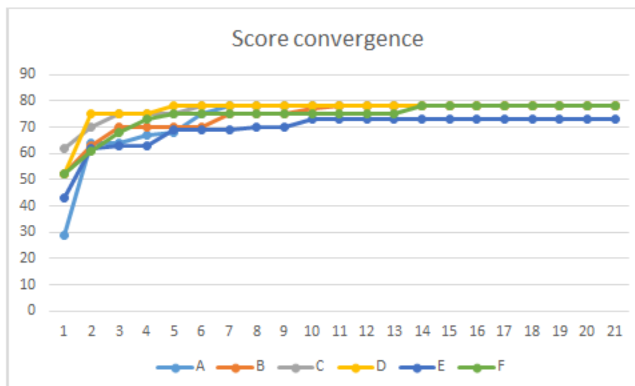

Fig. 2. Plot of convergence of the score during six different runs of the evolutionary process on the same input log from the simulation. Vertical axis shows maximum score for the population, horizontal axis shows number of generations considered to that point.

Figure 2 plots the maximum score at each generation for six different runs of the system on the same log. The graph shows several features of the system at work. The different values of the maximum score at the first generation is a result of the initial construction at random. The progressive increase of the scores shows the evolutionary algorithm at work. The rate of growth is slightly different across runs, consistent with the random nature of the evolutionary operators. Although there is a certain tendency to converge to the highest possible score at around the sixth generation, in some cases the curve does not altogether flatten out until the fourteenth generation. In view of this analysis we have considered that a value of 20 generations is adequate to ensure convergence to the maximum possible scores.

The procedure also needs to be tested on stories of more than two characters. The metric on character variety can be configured to assign maximum score to stories with between a range of values for the number of characters that take part in the story. This should allow the construction of stories of a higher number of characters. However, the logs being used as input did not include situations that imply elaborate three-way affective interactions of the kind picked out by the relevant metrics. The logs do include instances of subsets of events involving more than two characters that read reasonably well as a story. One such example is shown in Table 4

Ray proposes to Violet to weekend together as mates. Violet accepts Ray's invitation to weekend together as mates. Violet proposes to Ray to sleep together as mates. Ray accepts Violet's invitation to sleep together as mates. Violet proposes to Drew to weekend out as friends. Drew rejects Violet's invitation to weekend out as friends. Violet proposes to Drew to serious talk as friends. Drew accepts Violet's invitation to serious talk as friends. Violet proposes to Ray to invite dinner as mates. Ray rejects Violet's invitation to invite dinner as mates.

Table 4. Example of story draft obtained by configuring the character variety metric to assign maximum score to stories with between 2 and 4 characters.

This example shows how the story of the romantic relationship between Violet and Ray interweaves with the friendship between Violet and Drew. The difficulty for the proposed system to produce stories like this is that such a story does not involve the kind of significant changes in affinity between the characters that the metrics are designed to detect. The relationships between the three character in this story do not change significantly. Violet and Ray remains mates and Violet and Drew remain friends throughout. Furthermore, there is no indication in the story of what the affinities are between Ray and Drew and Drew and Ray. This makes this kind of story impossible to identify by the proposed procedure.

The difficulty that the system has in finding stories with more than two characters that exhibit the desired features is an indication that the proposed metrics, while valuable for selecting interesting interactions between two particular characters, may be too restrictive if they are to be applied to the complete set of possible affective threads arising from a story. Although stories that include information about all the possible directional affinities between all the characters involved would indeed be the densest possible instances of affective interaction between members of a collective, they would read out as exhaustive enumeration. Because of that, they would very likely give an impression of lack of focus. The stories that people generally like may mention a broad range of characters, but they tend to focus on a smaller subset–the protagonist and their main relations–rather than exhaustively cover them all. The metrics that we have proposed will need to be adapted to consider this aspect of stories, that they are not well suited to reflect in their current form.

Overall, the impression is that the stories are not very interesting. This is mostly due to the fact that the logs that are currently being used as input do not constitute very interesting simulations. The quality of the stories that can be sifted out of a story log is constrained by the interest of the events already present in the log in question. An obvious solution to this problem is to find more interesting simulations to process. This will be considered as future work. An alternative approach that might provide more informative data on the relative merit of the story sifting algorithm itself would be to develop an additional set of metrics capable of measuring the relative interest of the events in the whole log. That might then be used as a baseline in the sense that stories sifted from the log cannot add interest other than by intelligent selection.

In relation to the previous work cited in Section 2, the type of affinities considered in this system are very similar to those used by the Mexica system [17], in the sense that they are defined by a numerical value assigned to the directional affective connections between two characters. The affective threads proposed here as means of interpreting the evolution of the affinities between characters over a story can provide a basic computational approximation to concepts like character arcs [23]. The affective threads also have similarities with the emotional arcs proposed by Reagan [18]. However, emotional arcs differ in that they are not restricted to a specific pair of characters, but rather to the evolution of the overall affinity over the storyline. The features computed in the present paper over affective threads may be used to compute emotional arcs over stories.

The solution proposed here provides a computational procedure for sifting through simulation logs [20] in search for subsets of events that involve a set of characters undergoing significant transitions in their affinities to one another. It might also be employed as an ancillary procedure to identify which subset of a storywold might be more fruitfully selected to inform a process of composition of narrative discourse [4].

In constrast to the evolutionary solutions reviewed in Section 2.3 for building stories, the system presented in this paper does not actually build a story by combining independent fragments, but rather selects them from the set of events present in a given log. Nevertheless, this approach differs slightly from story sifting solutions described in Section 2.2 in the sense that, whereas story sifting solutions attempt to select existing units of action–say scenes in which a set of characters interact with a certain unity of space and time–the present solution considers sequences of events that appear in the log and involve the same set of characters, but which may actually have taken place at completely separate moments of time within the simulation. In that sense, this procedure does build a story by combining elements from a simulation log, with those elements possible being taken out of context.

5 Conclusions

The evolutionary approach to the task of sorting from the full set of events produced by a simulation those that have a potential for being an interesting story has shown a certain ability to identify subsets that exhibit desirable properties in terms of how the affinities between the characters involved develop. The interpretation of the sequence of the affective interactions in a given story in terms of a set of affective threads, the quantitative interpretation of these threads via a set of features, and the development of a set of metrics based on those features provide valuable tools for the development of fitness functions that can inform evolutionary procedures.

The proposed system allows identification of reasonable fragments of the life of the characters that show significant affective engagement between them. The fact that the resulting stories are not especially interesting is more a question of the type of simulation employed as input. We will consider exploring other possible simulations as input in future work.

Some of the potential shortcomings of the proposed procedure have been identified and discussed. One important issue is the assumption that significant affective interactions need to hold between all the characters involved. This assumption presents a requirement so strong that the system comes up empty when confronted with input logs that are not dense enough in affective interactions. Solutions to this problem will be developed as future work.

As future work, we want to consider the development of more elaborate metrics that identify occurrence of emotional arcs [18] over the affective threads identified for a story.

References

1. Behrooz, M., Swanson, R., Jhala, A.: Remember that time? telling interesting stories from past interactions. In: Interactive Storytelling: 8th International Conference on Interactive Digital Storytelling, ICIDS 2015, Copenhagen, Denmark, November 30-December 4, 2015, Proceedings 8. pp. 93–104. Springer (2015)
2. Concepción, E., Gervás, P., Méndez, G.: Ines: A reconstruction of the Charade storytelling system using the Afanasyev framework. In: Proc. of the Ninth International Conference on Computational Creativity, Salamanca, Spain. pp. 48–55 (2018)
3. Egan, K.: What is a plot? New Literary History **9**(3), 455–473 (1978)
4. Gervás, P.: Composing narrative discourse for stories of many characters: a case study over a chess game. Literary and Linguistic Computing **29**(4) (08/14 2014)
5. Gervás, P.: Storifying Observed Events: Could I Dress This Up as a Story? In: 5th AISB Symposium on Computational Creativity. AISB, AISB, University of Liverpool, UK (04/2018 2018)
6. Gervás, P.: Evolutionary stitching of plot units with character threads. In: WIVACE 2022 XVI International Workshop on Artificial Life and Evolutionary Computation (09/2022 2022)
7. Gervás, P.: Improving efficiency and coherence in evolutionary story generation. In: 14th International Conference on Computational Creativity. Waterloo, Ontario, Canada (06/2023 2023)
8. Gervás, P., Concepción, E., Méndez, G.: Evolutionary construction of stories that combine several plot lines. In: 11th International Conference, EvoMUSART 2022. Springer, Springer, Madrid, Spain (04/2022 2022)
9. Kreminski, M., Dickinson, M., Mateas, M.: Winnow: A domain-specific language for incremental story sifting. Proc. of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment **17**(1), 156–163 (Oct 2021)

10. Kreminski, M., Dickinson, M., Wardrip-Fruin, N.: Felt: A simple story sifter. In: Interactive Storytelling: 12th International Conference on Interactive Digital Storytelling, ICIDS 2019, Little Cottonwood Canyon, UT, USA, November 19–22, 2019, Proceedings. p. 267–281. Springer-Verlag, Berlin, Heidelberg (2019)
11. Kreminski, M., Dickinson, M., Wardrip-Fruin, N., Mateas, M.: Select the unexpected: A statistical heuristic for story sifting. In: International Conference on Interactive Digital Storytelling. pp. 292–308. Springer (2022)
12. Lee, O.J., Jung, J.J.: Story embedding: Learning distributed representations of stories based on character networks. Artificial Intelligence **281**, 103235 (2020)
13. Leong, W., Porteous, J., Thangarajah, J.: Automated story sifting using story arcs. In: Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems. p. 1669–1671. AAMAS '22, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (2022)
14. de Lima, E.S., Feijó, B., Furtado, A.L.: Procedural generation of quests for games using genetic algorithms and automated planning. In: 18th Brazilian Symposium on Computer Games and Digital Entertainment, SBGames 2019, Rio de Janeiro, Brazil, October 28-31, 2019. pp. 144–153. IEEE (2019). https://doi.org/10.1109/SBGames.2019.00028, https://doi.org/10.1109/SBGames.2019.00028
15. McIntyre, N., Lapata, M.: Plot induction and evolutionary search for story generation. In: Proc. of the 48th Annual Meeting of the Association for Computational Linguistics. pp. 1562–1572. Association for Computational Linguistics, Uppsala, Sweden (Jul 2010)
16. Méndez, G., Gervás, P., León, C.: On the Use of Character Affinities for Story Plot Generation, Advances in Intelligent Systems and Computing, vol. 416, chap. 15, pp. 211–225. Springer (02/2016 2016)
17. Pérez Ý Pérez, R., Sharples, M.: Mexica: A computer model of a cognitive account of creative writing. Journal of Experimental & Theoretical Artificial Intelligence **13**(2), 119–139 (2001)
18. Reagan, A.J., Mitchell, L., Kiley, D., Danforth, C.M., Dodds, P.S.: The emotional arcs of stories are dominated by six basic shapes. EPJ Data Science **5**(1), 1–12 (2016)
19. Rothlauf, F.: Representations for genetic and evolutionary algorithms. Springer (2006)
20. Ryan, J.: Curating Simulated Storyworlds. Ph.D. thesis, University of California Santa Cruz, CA, USa (12 2018)
21. Samuel, B., Summerville, A., Ryan, J., England, L.: A quantified analysis of bad news for story sifting interfaces. In: Interactive Storytelling: 14th International Conference on Interactive Digital Storytelling, ICIDS 2021, Tallinn, Estonia, December 7–10, 2021, Proceedings 14. pp. 142–156. Springer (2021)
22. de Silva Garza, A.G., y Pérez, R.P.: Towards evolutionary story generation. In: Colton, S., Ventura, D., Lavrac, N., Cook, M. (eds.) Proceedings of the Fifth International Conference on Computational Creativity, ICCC 2014, Ljubljana, Slovenia, June 10-13, 2014. pp. 332–335. computationalcreativity.net (2014)
23. Weiland, K.: Creating Character Arcs: The Masterful Author's Guide to Uniting Story Structure, Plot, and Character Development. Helping Writers Become Authors, PenForASword Publishing (2016), https://books.google.es/books?id=bMRgvgAACAAJ
24. Xie, Z.: Neural text generation: A practical guide. arXiv preprint arXiv:1711.09534 (2017)