Stories from Games: Content and Focalization Selection in Narrative Composition

Pablo Gervás

Universidad Complutense de Madrid,Spain,
pgervas@sip.ucm.es,
WWW home page: http://nil.fdi.ucm.es

Abstract. Game logs could be exploited as a source for reflection on user performance, with a view to improvement or simply the additional entertainment of reminescing. Yet in their raw form they are difficult to interpret, and sometimes only specific parts of the game are worth reviewing. The ability to produce textual narratives that rework these logs (or the interesting parts of them) as stories could open up this wealth of data for further use. This paper presents a model of the task of narrative composition as a set of operations that need to be carried out to obtain a linear sequence of event descriptions from a set of events that inspire the narration. As an indicative case study, an initial implementation of the model is applied to a chess game understood as a formalised set of events susceptible of story-like interpretations. Operating on simple representations of a chess game in algebraic notation, exploratory solutions for the tasks of content selection 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.

1 Introduction

The aftermath of a well-played game usually invites a process of reflection whereupon certain parts of the game are replayed in the mind (or on a blackboard or even on video). The purpose is sometimes to understand better a particular event in the game and sometimes simply to rejoice in it. This process can lead to improvements in performance, or it can be useful as a teaching aid. On a similar vein, digital entertainment generates a wealth of information in terms of traces of user activity in the form of game logs. These logs could be exploited as a source for reflection on user performance, with a view to improvement or simply the additional entertainment of reminescing. Yet in their raw form these logs are difficult to interpret, and reliving them in their original form would be too time consuming and not homogenously fruitful, as sometimes only specific parts of the game are worth replaying. The ability to produce textual narratives that rework these logs (or the interesting parts of them) as stories could open up this wealth of data for further use.

Automated composition of narratives from data is budding line of research within computational narratology, but elementary techniques are already available that can provide an initial approximation to this task [1–5]. The challenge to undertake is to provide a computational model of the traditional concept of *raconteur*: someone who can identify something that actually happened, leave out the boring bits, focus on the interesting ones, and tell it in a way that makes it entertaining. The task itself can be broken down into several stages:

1. he selects the fragment of real life he wants to tell
2. he filters out the non-relevant parts to obtain his initial material
3. he imposes a certain order on the events that he wants to tell (from here on he has a narration sequence)
4. he decides where in that sequence to insert the descriptions he intends to use from his initial material (he now has a sequence that switches between narration and description)
5. he decides how much of that sequence he can assume that his listeners will infer without effort from the rest
6. he produces the right form for the remaining sequence

Steps 3 and 4 may be somewhat mixed or carried out in reverse order, depending on the style of the raconteur. Step 5 may be ommited. Step 6 involves linguistic skills that are a huge problem in themselves.

The present paper attempts to address these challenges from an engineering point of view: given an exhaustive record of all moves made in a given game, find a way of telling what happened as a linear sequence that may be later converted into text, trying to maximize coverage, minimize redundancy, and achieve a certain natural fluency. Chess has been chosen as an initial case study because it 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. Yet it also allows very elementary interpretations of game situations in terms of human concepts such as danger, conflict, death, survival, victory or defeat, which can be seen as interesting building blocks for story construction.

The paper reviews existing models related to the narrative composition task, describes the proposed computational model, presents the case study for narration of chess games and finishes with discussion and conclusions.

2 Previous Work

A number of models of related tasks and elements arising from different fields of research are reviewed in this section to provide background material for the discussion. Due to the breadth of fields under consideration, exhaustive review in any one of them is beyond the scope of the paper. An attempt has been made in each case to gather here the set of elementary concepts in each field that are relevant for the understanding of the arguments in the paper.

2.1 Narratology

According to many theorists, narrative has two components: what is told (what narrative is: its content, consisting of events, actions, time and location), and the way it is told (how the narrative is told: arrangement, emphasis / de-emphasis, magnification / diminution, of any of the elements of the content). These have been named different ways by different researchers, story and discourse, *histoire* and *discours*, *fabula* and *sujzet*. There are alternative analyses that postulate different subdivions. Even between theories that agree on having just two levels of analysis there seem to be many subtleties that cast doubt on whether the same thing is meant by the different words. This presents a serious obstacle for researchers from the computational field trying to address the treatment of stories in any form. In order to avoid ambiguity, we will restrict our analysis here to three levels of conceptual representation of a story, and refer to these as the *story* (the complete set of what could be told, organised in chronological order of occurrence), the *plot* (what has been chosen to tell, organised in the order in which it is to be told) and the *narrative* (the actual way of telling it).

Narratologists, who specialize in the study of narrative, consider the concept of *focalization* [6] as the way in which a narrator restricts what he is telling about a particular scene to what might have been perceived by someone present in that scene. This may be one of the characters if the scene is told in the first person, or the narrator himself as if he had been present (if the story is told in the third person). This has an interesting implication in the fact that, through focalization, narrative discourse (and thereby the structure of stories) is influenced by the perception of space: events that take place simultaneously in different locations that cannot be perceived at the same time (this may be different cities but also different neighbouring rooms separated by a wall) usually require different narrative threads.

2.2 Natural Language Generation

The general process of text generation takes place in several stages, during which the conceptual input is progressively refined by adding information that will shape the final text [7]. During the initial stages the concepts and messages that will appear in the final content are decided (*content determination*) and these messages are organised into a specific order and structure (*discourse planning*), and particular ways of describing each concept where it appears in the discourse plan are selected (*referring expression generation*). This results in a version of the discourse plan where the contents, the structure of the discourse, and the level of detail of each concept are already fixed. The *lexicalization* stage that follows decides which specific words and phrases should be chosen to express the domain concepts and relations which appear in the messages. A final stage of *surface realization* assembles all the relevant pieces into linguistically and typographically correct text. These tasks can be grouped into three separate sets: *content planning*, involving the first two, *sentence planning*, involving the second two, and surface realization. An additional task of *aggregation* is considered, that involves merging structurally or conceptually related information into more compact representations ("Tom fled. Bill fled." to "Tom and Bill fled." or "the boy and the girl" to "the children"). Aggregation may take place at different levels of the pipeline depending on its nature.

3 A Computational Model

The type of narrative that we want to address in this paper involves a linear sequential discourse where only a single event can be told at any given point. Yet reality is not like that. Events to be reported may have happened simultaneously in physically separated locations, and constitute more of a cloud than a linear sequence, a volume characterised by 4 dimensional space time coordinates. Composing a narrative for such an input involves drawing a number of linear pathways through that volume, and then combining these linear pathways (or a selection thereof) together into a single linear discourse. This type of linear pathway is sometimes referred to as a *narrative thread*. From a narratological point of view, this task can be related to the identification of appropriate focalization decisions for conveying a given material. Focalization, understood as the decision of which character the narration should follow, and how much of the environment around him at each point should be conveyed to the reader of the narrative, divides the perception of reality into individual fibres (one for each possible focalizer character) that are linear and sequential in nature. For each character involved in the set of events to be conveyed, a possible focalization fibre can be drawn.

To provide a preliminary benchmark for the various intuitions outlined in the rest of the paper the simplest approximation to a case study that could be conceived is described in this section. Characters will be chess pieces. To model the way humans tell stories in terms of narrative threads, based on subsets of the events that can be said to be experienced by a particular character, each piece is assigned a field of perception corresponding to a N x N square around the position of the board in which it is standing. The value of N is considered the *range of perception* of that particular piece. For a given piece with range of perception N, all possible partitions of the chess board in N x N squares would be possible locations in our world model - depending on the positions of characters on the board. The set of events that we would be considering is the set of all moves involved in a given game.

A basic software implementation has been written that reads a description of a chess game in algebraic notation (see Table 1) and builds for it a representation in terms of threads focalized on a given character assuming a certain range of perception around him.

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.

1. e4 c5	16. Bxe2 Be6
2. Nf3 d6	17. Rfd1 Rfd8
3. d4 cxd4	18. Bc5 Rd5
4. Nxd4 Nf6	19. b4 a5
5. Nc3 g6	20. Bf3 Rxd1+
6. Be2 Bg7	21. Rxd1 e4
7. Be3 O-O	22. Bxe4 Bxc3
8. O-O Nc6	23. Bxc6 Rc8
9. h3 d5	24. b5 Bxa2
10. exd5 Nxd5	25. Bd4 Bb4
11. Nxd5 Qxd5	26. Be5 Be6
12. Bf3 Qc4	27. b6 Rxc6
13. Nxc6 bxc6	28. b7 Rb6
14. c3 e5	29. Rd8+
15. Qe2 Qxe2	1-0

Table 1. Algebraic notation for an example chess game

3.1 Representing the Data

An *event* is something that happens with a potential for being relevant to a story. Events occur at a specific location at a given moment of time. They may have preconditions and postconditions. In order to have a generic representation, each event is considered to have an associated *event description* that allows both a descriptive and a narrative component. The descriptive component of an event description contains predicates that describe relevant preconditions. The narrative component of an event describes the action of the event itself, but it may also contain additional narrative predicates describing the effect of the action. For instance, if someone moves from one place to another, elements at the original location may be left behind, and elements in the final location appear. These are not considered separate events but included in the description of the moving event.

A *fibre* is a sequence of events that either involve or are seen by a given character. It represents a focalized perception of the world. The extent of information that is included in the events for a given fibre is defined by the range of perception that is being considered and by the presence of any obstacles to perception in the surrounding environment. It may also be affected by the direction in which the focalizer is facing, or where he is focusing his attention. As these further refinements require advanced authorial decisions, we decide that the standard representation should include all the information within perceptive range, and leave the decision of whether to mention it to a later stage of composition.

The task of *heckling*[1] involves establishing the range of perception, tracking the set of all possible characters involved in the events to be narrated, and

[1] From the process of extracting a set of fibres from flax or fleece, to be later spun into yarns.

for each character constructing a fibre representation that includes descriptions of all the events that the character initiates, suffers or perceives. These event descriptions will appear in the fibre in chronological order.

3.2 A Computational Procedure

From a given game, a large number of fibres, longer or shorter, and at different ranges of perception can be produced. Most of them overlap one another, and most of them will contain descriptions of events that are not interesting in themselves and not relevant to the overall flow of the game. To select among them a selection of subsets that might constitute acceptable tellings of the game we have considered an evolutionary approach. This allows us to concentrate our effort on the definition of fitness functions that capture desireable qualities of a good story, both at the domain level and in terms of the structural properties it should have. We therefore need to define an initial population of fibres, and a set of operators for crossing and mutating them, using these fitness functions to rate successive generations. This would constitute an evolutionary approach [8] to the task of content and focalization selection.

The concept of a *draft* that holds the current best solution for a given telling of the game under consideration and which gets progressively modified towards an optimal solution, is fundamental to the proposed model. The concept of *reviser*, a module that operates on a draft to progressively improve it, captures this concept of progressive modification. The model will operate not on a single draft but over a *population* of candidate drafts.

We allow a set of alternatives for the creation of the drafts in the initial population. To this end we introduce the concept of *babbler*, a module in charge of producing an initial draft. By allowing a population of babblers to produce the initial population, we introduce the possibility of relying on more than one technology to produce them. Solutions for babblers that have been explored so far generate a set of single fibre drafts which are then combined into more complex fibres by subsequent stages of the system.

Drafts need to be evaluated for conformance with the desired qualities for the telling of the game, and the results of this evaluation need to be taken into account in any subsequent operations on the draft. The concept of a *judge*, a module capable of evaluating partial results according to desired criteria, captures this idea. The various judges assign scores on specific parameters (designed to isolate structural properties such as redundancy, coverage, cohesion or overall interest):

- *uniqueness*, measures the numbers of events made or seen by the focalizers of the fibres involved in a telling that are not already covered by some other fibre in the telling, as a percentage of the total number of narrative predicates in the yarn (which should correspond to all the events mentioned),
- *density* measures the ratio between captures and narrative predicates (the number of narrative predicates that are captures as opposed to just moves),
- *coverage*, measuring the percentage of moves of the game that are actually covered by a given telling,

- *cohesion*, measuring the ratio between the number of moves covered by the longest fibre in a telling and the number of moves covered by the telling (not necessarily the game, as parts of it may have been lost, and that is measured elsewhere by the coverage metric),
- *number of focalizers*, assigning high scores to tellings that have a number of different focalizer within a specified range (currently set between 2 and 6).

The definition of density is based on the assumption that captures constitute more emotionally charged events than other moves. An overall score for each draft is obtained by averaging the value of all individual scores received by the draft. Weighted combinations may be introduced at later stages if they are seen to improve the quality of the final results.

Evolutionary approaches rely on a set of cross over and mutation operators [8]. In our framework, cross over functionality is introduced by *mixers* and mutation functionality by *revisers*.

Two types of mixers are considered:

- one that given two different tellings, generates another two in which fibres ocurring for the same focalizer in both, but with different coverage of the actual set of events it took part in (either as an active participant or as an observer),[2] have been swapped
- one that given two different tellings produces a single telling that combines all the significantly different[3] fibres found in both

During one generation, each mixer is applied as many times as there are drafts in the population, and each time it combines two drafts selected at random.

Revisers rely on scores assigned by judges to introduce changes to drafts. Reviser act as mutator operators, taking an existing draft and producing a new draft that differs from the original in some way. The set of revisers includes:

- dropping certain fibres from a telling (for each available draft, produce a new draft by dropping fibres at random with a likelihood of 10 %)
- trimming certain fibres from a telling (for each available draft, produce a new draft with the same number of fibres but dropping event descriptions from the original fibre at random with a likelihood of 5 %)

During one generation, each reviser is applied to all the drafts in the population.

At the end of each generation, the size of the population has multiplied significantly. The resulting set of drafts are scored, and only top scoring ones are allowed onto the next generation. The number of drafts allowed to survive from each generation is a configuration parameter currently determined empirically.

[2] Differences in coverage may arise from the fibre having been trimmed as a result of revision or by virtue of having been originally constructed based on a different value for the range of perception.

[3] Two fibres are considered significantly different unless the set of events covered by one subsumes the set of events covered by the other.

3.3 Results

Running the system produces a population of drafts that constitute possible alternative ways of telling the story by focusing on one or another set of characters. With a view to explainig the intuitions behind the described evolutionary set up, an example of system run with a simple configuration is described below. These tellings are produced by a combination of a babbler that produces a different draft for each available fibre, a mixer that combines drafts by merging all their constituent fibres, and the full set of judges described.

For instance, for the game shown in Table 1, a number of possible tellings are summarised in Table 2. The top scoring single fibre candidate is draft 6, telling the story of the game that focuses on the second white pawn using a range of perception of 3. This is the top scoring draft on coverage (78) that has a single fibre. This pawn sees much of what happens in the game. It also has a top score on uniqueness (100), because no one else is reporting the same things. The overall top performer is draft 114, that tells the game through the eyes of the fifth black pawn and the left white bishop. The pawn uses a smaller range of perception (2) than the bishop (3). They achieve a top score on coverage of 96, but get penalised on uniqueness (73), because there is a large number of events that are reported by both of them (in bold in their list of moves). For completeness, data for two additional drafts (numbers 97 and 124) are shown, both with an overall score of 70 but using two and three focalizers respectively. Maximum coverage (98) is achieved with three focalizers in 124, but uniqueness drops down to 53, bringing the overall score down. This example shows how the metrics of the different judges interact to achieve a balance that leads to successful tellings of the game.

4 Discussion

The proposed system constitutes a computational implementation of two tasks relevant to narrative compostion: selection of relevant material by filtering out less relevant information (corresponding to the content determination subtask in natural language generation pipelines) and determination of appropriate focalization choices (a subject studied by narratology, though usually considered only during analysis of literary texts in terms of how focalization has been established by human authors). From the point of view of natural language generation, the task of content determination for narrative had not to our knowledge been explored computationally before. From the point of view of narratological theory, the model as described captures the concept of focalization as a computational generative task, and relates it closely to the input data, the decision criteria, and the computational processes being modelled.

4.1 Comparison with Other Narrative Composition Systems

A number of related efforts exist to automatically derive narratives from sport games [2],[3],[4]. These efforts operate on input data in the form of statistics on

Dn	Fc	Un	Cv	Ch	Dn	Os	# F	# m	Piece	Pr	List of moves
6	50	100	78	100	14	68	1	45	wp2	3	1, 2, 5, 6, 7, 9, 11, 13, 15, 18, 19, 20, 21, 22, 23, 24, 25, 27, 28, 29, 30, 31, 33, 35, 36, 37, 38, 39, 40, 41, 42, 43, 44, 45, 47, 48, 49, 50, 51, 52, 53, 54, 55, 56, 57
114	100	73	96	89	10	73	2	55	bp5	2	2, 4, **6**, 8, 10, 12, 14, **16**, **18**, **19**, **20**, **21**, **22**, **24**, **25**, **26**, **28**, **30**, **31**, **32**, **35**, **36**, **39**, **40**, **42**, **43**
									lwb	3	1, 3, 5, **6**, 7, 9, 11, 13, 15, **16**, 17, **18**, **19**, **20**, **21**, **22**, 23, **24**, **25**, **26**, 27, **28**, 29, **30**, **31**, **32**, 33, **35**, **36**, 37, 38, **39**, **40**, **42**, **43**, 44, 45, 46, 47, 48, 49, 50, 51, 52, 53, 54, 55, 56, 57
97	100	73	92	73	14	70	2	53	wp7	3	1, 3, 5, 6, 7, 11, 13, 15, 17, 18, 19, 20, 21, 22, 23, 24, 25, 28, 29, 30, 31, 33, 35, 36, 39, 40, 41, 42, 43, 45, 49, 51, 57
									lbb	3	2, 4, 6, 8, 12, 14, 16, 18, 19, 20, 21, 22, 24, 25, 26, 28, 32, 34, 35, 36, 37, 39, 40, 42, 43, 44, 45, 46, 47, 48, 49, 50, 51, 52, 53, 54, 55, 56, 57
124	100	53	98	87	13	70	3	56	wp7	3	1, 3, 5, 6, 7, 11, 13, 15, 17, 18, 19, 20, 21, 22, 23, 24, 25, 28, 29, 30, 31, 33, 35, 36, 39, 40, 41, 42, 43, 45, 49, 51, 57
									bp6	2	4, 8, 10, 12, 14, 18, 19, 20, 21, 22, 24, 28, 32, 34, 36, 40, 42, 44, 48, 51, 52, 57
									lwb	3	1, 3, 5, 6, 7, 9, 11, 13, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 32, 33, 35, 36, 37, 38, 39, 40, 42, 43, 44, 45, 46, 47, 48, 49, 50, 51, 52, 53, 54, 55, 56, 57

Table 2. Examples of tellings: draft number (Dn), number of fibres (# F), number of moves (# m), range of perception of the fibre (Pr), focalizer name (Piece), and scores for focalizer count (Fc), uniqueness (Un), coverage (Cv), cohesion (Ch), density (Dn) and overall score (Os).

a given game, and produce texts in the manner of newspaper articles covering similar games. These efforts, for instance, are heavily mediated by the set of data they start from, which arises from purpose specific abstraction and filtering of the set of real world events, driven by the purpose of the desired story. Additionally, the choice of news article as target form imposes stylistic choices in terms of focalization (the stories are not focalized on specific characters) and the rhetorical structure (for instance the need to provide a summary of the game at the beginning, then review the highlights, then comment on any relevant details). The input data considered in these cases already involve a severe operation of abstraction, in the sense that they do not constitute raw perceptions of reality but purpose-specific abstractions such as statistics on percentage of succesful attempts, relative number of fouls between opposing teams, or total time that a team had control of the ball. In such cases, a large portion of the effort of converting the real life data into story material has already been achieved by the stage of drawing statistics from the game. As such, they differ greatly from the type of input data and output text being considered in this paper. For the model described here, input data would be a complete log of all events that took place in the game, and the output text would be closer to a multi-character novel describing the game, or the script of a film telling the story of the game.

The task of narrative composition has also been explicitly addressed by Hassan et al [1] and Gervás [5]. Hassan et al addressed the task of generating stories from the logs of a social simulation system. Gervás focused on the task of generating a natural language rendering of a story extracted from a chess game for a given set of pieces and with focalization set at a prefixed range of perception (squares of size 3 were used). An analysis of the differences between these systems is best carried out in terms of the 6 stages described in section 1 for the task of a *raconteur*.

Of those stages, step 1 would correspond to selecting a particular game, or particular social simulation, and in all cases is already decided in the input received by the systems. The system described in the present paper models step 2, corresponding to filtering out the non-relevant parts. This step was not addressed by Hassan et al [1]. The procedure employed in the present case relies heavily on the fitness functions implemented in the judges to establish the filtering criteria, though it is the mixers and the revisers that carry out the executive work of tailoring the drafts. Step 3 (imposing a particular order on events to be told) is addressed by all three systems, with Gervás [5] being the one that devotes most effort to it. The present system relies on the solution for step 3 developed in Gervás [5]. This step presents interesting features in the context of the case study on chess, and is addressed below under 4.2. Step 4 (generating and inserting descriptions at appropriate places in the narrative sequence) has so far only been addressed in Gervás [5]. Step 5 (establishing how much of the intended content can be omitted and left for the reader to infer) requires a rich model of what the reader knows and has so far been sidestepped by all previous research efforts. Step 6 (rendering the resulting conceptual discourse as natural language) was addressed by Hassan et al [1] in terms of template-based NLG and by Gervás [5] with a more refined NLG system that included referring expression generation, lexical and syntactic choice, and grammar-based surface realization.

4.2 Generalizing beyond Chess

Chess games present the advantage of having very clear temporal and spatial constraints, and constituting at heart a sketchy representation of one of the most dramatic settings for human experience: war. In that sense, it provides a good ground for simple experiments, and it is not intended as a contribution in itself but as an illustrative example of the operation of the model of sufficient simplicity to be describable within the size restrictions of a paper such as this. Two aspects are identified as problematic with the use of chess games as a case study, one related to the tasks of content selection and focalization choice and one more related to the task of establishing a sequential order on the narrative (corresponding to step 3 mentioned in section 1).

First, adequate representation of opportunities and threats in a chess game involves some serious representational challenges [9]. For more elaborate experiments, the fitness functions employed would have to be refined to better capture this complexity. Although significant progress may be made on this point, the effort invested is unlikely to lead to compelling narratives or narratives that bring insight on the task of narrative composition. It might be more fruitful to identify richer domains for further case studies, and to apply the engineering effort to whatever elements are considered of narrative interest in them.

Second, the chronological structure of a chess game is in truth purely sequential, in contrast with the sets of events that would be considered when narrating from real life events. This has not been considered a serious obstacle in as much as focalization breaks up the set of events into separate views corresponding to different fibres, and the numbering or moves in the game provides a good indication of relative chronology of events within any given fibre and across fibres. For these reasons, it is considered advisable to explore further investigation of the suitability of the model when applied to case studies in other domains that are richer in terms of their representation of time and space and that may lead to more compelling narratives with a stronger human interest. The search for alternative domains is made difficult by the need to obtain for them reliable records of all events, both relevant and irrelevant to the story that may be told. Only if this condition is satisfied can the corresponding problem be considered equivalent to the human task that we want to model.

These two points suggest strongly that logs of video games would be a very good domain of application, as they satisfy all the described requirements. Situations and events of interest can be established for each particular game, and means would have to be found for identifying them computationally. Existing research on plan recognition [10–12] may be of use in addressing this task. Logging procedures can be defined for the chosen game, to ensure that game events, both relevant and irrelevant to the story that may be told, get registered.

5 Conclusions

The evolutionary solution presented in this paper constitutes a valuable contribution to the growing corpus of research on narrative composition. The tasks of content determination and focalization choice in narrative had not been addressed computationally before. The evolutionary approach provides a breakdown into subtasks that has lead to interesting insights in terms of specific knowledge-based operations that need to be carried out during composition. These operations provide a first computational approximation to decision making about focalization, as opposed to analysis of how focalization has been configured. These tasks can also be confortably placed within generally accepted task divisions for natural language generation. The preliminary implementation for the particular case of a chess game has shown that a set of possible sequential tellings with acceptable structural properties can be obtained from a formal log of considerable complexity. This opens the possibilities for addressing richer case studies for more complex games.

References

1. Hassan, S., León, C., Gervás, P., Hervás, R.: A computer model that generates biography-like narratives. In: International Joint Workshop on Computational Creativity. London. (2007)
2. Allen, N.D., Templon, J.R., McNally, P., Birnbaum, L., K.Hammond: Statsmonkey: A data-driven sports narrative writer. In: Computational Models of Narrative: AAAI Fall Symposium 2010. (2010)
3. Lareau, F., Dras, M., Dale, R.: Detecting interesting event sequences for sports reporting. In: Proc. ENLG 2011. (2011) 200–205
4. Bouayad-Agha, N., G.Casamayor, Wanner, L.: Content selection from an ontology-based knowledge base for the generation of football summaries. In: Proc. ENLG 2011. (2011) 72–81
5. Gervás, P.: From the fleece of fact to narrative yarns: a computational model of narrative composition. In: Proc. Workshop on Computational Models of Narrative 2012. (2012)
6. Genette, G.: Narrative discourse : an essay in method. Cornell University Press (1980)
7. Reiter, E., Dale, R.: Building Natural Language Generation Systems. Cambridge University Press (2000)
8. Goldberg, D.E.: Genetic Algorithms in Search, Optimization and Machine Learning. 1st edn. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA (1989)
9. Collins, G., Birnbaum, L., Krulwich, B., Freed, M.: Plan debugging in an intentional system. In: Proc. IJCAI 1991. (1991) 353–359
10. Fagan, M., Cunningham, P.: Case-based plan recognition in computer games. In: Proc. of the Fifth ICCBR, Springer (2003) 161–170
11. Ramirez, M., Geffner, H.: Plan recognition as planning. In: Proc. of IJCAI-09. (2009)
12. Synnaeve, G., Bessière, P.: A Bayesian Model for Plan Recognition in RTS Games applied to StarCraft. In: Proc. of AIIDE 2011, Springer (2011) 79–84