The AIIDE-17 Workshop on
Intelligent Narrative Technologies
WS-17-20

Towards a Computational Measure of Plot Tellability

Leonid Berov
University of Osnabrück
49076 Osnabrück, Germany
lberov@uos.de

Abstract

Measuring the quality of plot is a desirable feature for computational narrative systems. One of the notions of plot quality used in narrative theory is called tellability, which can be derived from certain structural properties, namely the types of events present and the way they are connected. These structures include not only actualized events, but also take into account virtual plans and the affective valencies of events. The present paper introduces Marie-Laure Ryan's tellability principles and suggests to computationally model them using an affective multi-agent simulation system. It discusses how such an approach implies a broader understanding of plot than commonly assumed and analysis several existing narrative systems under these considerations. Furthermore, it introduces a plot-graph formalism that allows the computational representation and analysis of the extended plot understanding. An approach to automatically generating the plot-graph is suggested in the context of the introduced multi-agent simulation system.

Introduction

Like for all computationally creative systems, automatic quality assessment is a long-standing goal of computational storytelling systems, since it enables a feedback loop that helps the machine to improve, curate or even comment on its creations (Gervás 2009). But even systems that are not interested in appearing creative can benefit from a quality measure, since it allows to make informed decisions during the generation process itself.

The quality of a narrative is one of its features and by that merit the object of study of narrative poetics. It is possible to differentiate two questions: "what makes a story worth telling" and "how to tell a story well" (Ryan 1991, p. 149). The former tries to establish the quality of plot, while the latter addresses the quality of its rendering, the discourse and text. The present paper will be concerned only with plot.

By taking a prescriptive stance the poetics of plot can address another set of questions: "what is a good plot", "what is a bad plot" or also "what is no plot at all"? From this perspective it becomes apparent that an understanding of plot, as well as its computational representations, should take into consideration narrative phenomena used to assess tellability.

The present paper will introduce one specific approach to plot quality, Marie-Laure Ryan's (1991) *tellability*. Its discussion will make apparent that a simple understanding of plot as a (chronologically or causally) ordered set of actions is not far-reaching enough and that other, virtual, components should also be included into conceptions of plot. Interestingly enough, this is an understanding that has been recently more and more adopted by several computational systems (Ware and Young 2014; Chang and Soo 2008; Pizzi and Cavazza 2007), but due to differing considerations.

After outlining Ryan's understanding of plot quality the paper will propose how it might be computationally modeled. It will do so in the context of a multi-agent simulation system, and by deriving the introduced tellability principles from execution properties of the agent architecture. In order to do so it will also introduce a plot-graph formalism that allows a functional analysis of plot units.

The aim of this paper is to open a discussion on means of implementing the proposed computational model and on ways of synthesising the qualitative principles imported from narratology into a quantitative measure that can be used for evaluation by computational storytelling systems.

Tellability and Embedded Narratives

Tellability is a concept used by narratologists to describe the potential suitability of a configuration of events to be rendered in a story (Labov 1972; Prince 2003; Abbott 2014). Because of its pre-textual nature Ryan (1991, p. 149) uses it as a quality measure for plot, independent from its possible renderings on a discourse or textual level. In her framework she distinguishes several aspects contributing to tellability. The external aspect considers facts like the intended audience or the context of the storytelling occasion: plots enjoyed by infants, for instance, rarely interest mature readers. The (internal) substantial aspect is related to the semantic meaning of the involved events. As example Ryan (p. 154) presents a whole catalogue of topics and motifs of interest, while Bruner (1991, p. 11) summarizes that the events should report the breaching, violation or deviation from a canonical script (in the Shankian sense). The (internal) formal aspect of tellability is independent of content and context and is concerned only with structural properties: with ways in which types of events can be combined in order to be tellable. It appears that this aspect is the most likely to generalize well for computational storytelling systems, which may differ in context, audience or the intended genre but are generally all concerned with exploring ways to connect different events. Ryan (1991, pp. 155, 249) identifies four of these intrinsically aesthetic plot structures, which are summarised below. Personally, I also consider a part of Ryan's (p. 153) *dynamic* aspect to describe rather a structural property and for that reason append it below. Thus, each instance of the following principles is argued to increase a plots tellability:

1. *semantic opposition*: "reversals in the fortunes of characters" and contrasts between "goals of characters with the results of their actions",
2. *semantic symmetry*: structural similarities in sequences of events pertaining to different characters, or to the same character at different stages of the plot,
3. *functional polyvalence*: the same event fulfilling several narrative functions at the same time,
4. *suspense*: a delay between the adoption of a goal by a character and its achievement/failure,
5. *dynamic points*: the violation of a character's expectation.

However, the theory (p. 251) admits the caveat that an unbounded proliferation of each of these principles might result in repeating, predictable structures and would thus be detrimental to tellability. It is proposed that an aesthetic balance between the individual principles and with the overall structure must be maintained. Unfortunately, Ryan leaves it open how to detect a breach of this balance.

Before embarking on a discussion of how tellability principles can be formulated computationally, important remarks need to be made on the underlining understanding of plot. Observing that a necessary condition for tellability is the existence of a plot, Ryan (p. 154) argues that a plot consists of at least one conflict and one or more attempts at solving it. Such conflicts may exist between the beliefs, wishes or obligations of a character within herself, with other characters or the actual state of the story world. She observes that (within reasonable bounds) the tellability of an attempted resolution of a conflict increases with its complexity: The story of a hero who plans to conquer the heart of a princess, carries out the plan and succeeds is far less interesting than one about a hero who tries to execute his plan, fails due to force majeure and has to improvise to succeed.[1] However, this means that in order to analyse tellability we need to take into account not only the events that actually take place in a narrative (usually actions), but also plans that fail partially or even remain completely unattempted (that is, actions untaken). This amounts to an extension of the concept of plot to mean "[...] a complex network of relations between the factual and the nonfactual, the actual and the virtual" (Ryan 2013). Under this interpretation a more complex network represents a more tellable plot.

This virtuality resides in what Ryan calls *embedded narratives*[2]: story-like constructs containing characters' subjective representations of past or future states of the story world. These constructs contain not only actualized or unactualized plans, but also characters affective appraisal of past events, beliefs about the (past and present) states of the story world or desires about its future states. "The aesthetic appeal of a plot is a function of the richness and variety of the domain of the virtual, as it is surveyed and made accessible by those private embedded narratives" (Ryan 1991, p. 156). The scope of the embedded narrative concept is so big that narratologist Alan Palmer interprets it to represent the whole functioning of the fictional mind (Palmer 2004, p. 121f). It is interesting to note that in a computational model the 'functioning of the fictional mind' could be interpreted to mean the runtime behavior of an instance of the character model. This observation—that the internals of program execution can be used for quality assessment—has also been made from a computational perspective by Gabriel (2016), who uses logs and traces of his InkWell system to discuss the quality of generated haiku. The next section will explore this idea further.

With this extended understanding of plot in place it is possible to return to the discussion of tellability. For the principles introduced above this crucially means that their detection can rely not only on actualized actions, but should also take into account all parts of characters' embedded narratives. It also allows Ryan (1991, p. 228) to formulate a *rule of narrative closure* that establishes when a plot is well formed, and which can act as a post-condition in addition to the pre-condition already introduced above: A plot is well formed, when all plans that were initiated are also concluded, that is when all character goals have either succeeded or conclusively failed.

Towards an Implementation

In the following, a computational model is suggested that outlines how the above notion of embedded narratives can be implemented and how this can be used to identify instances of the introduced tellability principles. In its suggested shape the model constitutes only a representational formalism for narratives, however, I see no fundamental obstacles that would prevent computational or interactive storytellers from being built on top of it.

Character Architecture As I have suggested in (Berov 2017), a multi-agent simulation system based on the Belief-Desire-Intention (BDI) architecture (Rao and Georgeff 1995) extended with personality-based affective reasoning can be used in order to model Ryan's (1991) possible-worlds conception of fictional character. In such a system, each agent maintains its own propositional knowledge base (its *beliefs*) and a set of objectives (*desires*). It then uses procedural reasoning—whereby selecting applicable (partial)

[1] See Ryan (1991, p. 157ff) for examples from narratives like "The Crow and the Fox", "Cinderella" or "The Decameron".

[2] In the games literature embedded narratives typically refer to a mode of communicating a story embedded in the gameplay (Wei 2010). In this paper the use is different in that the embedded story is not necessarily directly conveyed to the consumer, but can be reconstructed from a combination of character actions and thought report.

plans is interleaved with perception and action-execution—to attain its *intentions*, which are a non-conflicting subset of the desires. The separation of agent-internal states makes BDI a promising character model for our purpose because it captures the subjective nature of embedded narratives: the same actual events can result in different beliefs, intentions or adopted plans for different characters. The system can also be used to create the affective part of embedded narratives. Emotions are primarily generated by the agent's reasoning cycle as a means of short-term affective appraisal of external and internal events. All arising emotions are aggregated into a mid-term affective state vector called mood, which affects the agents planning. Since emotions arise as an effect of past events, are preserved in the mood over time, and affect the agents behavior, they can also be used in our context as a part of the agents embedded narrative. Thus, the above system is capable of capturing retrospective embedded narratives as beliefs and emotions, while prospective embedded narrative are captured as intentions and partial plans. Coming back to Palmer's interpretation of embedded narratives as the functioning of the fictional mind, we can observe that it also holds true in the suggested BDI model. Embedded narratives are represented through *internal events* from the execution of individual agents' reasoning cycle: the addition or removal of beliefs, the selection or termination of intentions, the procedural selection of partial plans or plan failures, and the arising of emotions.

Conflict and Plot As mentioned above, a pre-condition of tellability is the presence of conflict and at least one solution attempt. Using the introduced agent architecture it becomes possible to capture the presence of conflict in a plot as follows:

- Each adopted intention represents an actual conflict between a character's wishes and the story world.
  - If the resolution of two agents' intentions would result in incompatible state changes of the environment, then this additionally represents an inter-character conflict.
  - The resolution of an agent's intention can also result in a state of the environment that is incompatible with this agent's desires, which additionally establishes an intra-character conflict. While such conflicts may not necessarily influence the actions actually performed by a character, they are never the less part of the virtual and by that virtue relevant for plot.
- An epistemic conflict occurs when an agents has a false belief and this belief is used during planning at least once.
- If the resolution of an agent's desire or intention would result in an environment that is incompatible with another agent's desire, but the later desire is never selected as intention, then this establishes a latent conflict that is again part of the virtual.

An intentional solution for a conflict in this framework can only occur in the form of a plan. This means that the tellability pre-condition can be fulfilled only by plots that have at least one actual conflict. The virtual and epistemic conflicts that can accompany an actual conflict can be seen as vital ways to increase tellability because they increase a plots complexity. Of special interest are failing plans because they indicate the presence of an embedded conflict: either the beliefs used for planing were inaccurate, or another agent's conflicting intentions caused a plan to fail.

The post-condition of narrative closure can be captured in equal terms. An execution-run of a multi-agent simulation represents a well-formed plot at each step at which the following holds: (1) the tellability pre-condition holds, and (2) no agent has an active intention. It is conceivable that due to the self-perpetuating nature of multi-agent systems no single step can be identified at which the second condition holds. In such a case a relaxed condition (2') can be used that allows to select for each agent an individual step as end step iff at this step the agent's intention stack is empty. This allows to represent open-ended plots, in the sense that at the selected end of the plot individual character's reactions are possible but not included.

Tellability Principles Having established that it is possible to capture the notion of a well-formed plot, it remains to also represent the formal principles of tellability in terms of the character architecture. This makes it necessary to already make one reference to the hitherto unmentioned concept of functional units, which will be introduced in the next section. The tellability principles can be represented as follows:

1. *semantic opposition*:
   (a) a reversal in a character's fortune occurs each time two subsequent emotions differ in valence sign[3] (pos/neg),
   (b) a contrast between a character's goals and it's achieved results occurs each time a plan failure is detected and accompanied with an emotion of negative valence,
2. *semantic symmetry*: structural similarities occur each time the same sequence of internal events (or valencies, or functional units, depending on the employed level of abstraction) appears in different characters, or within the same character at different steps of the simulation,
3. *functional poly-valence*: an event fulfills several narrative functions if it is part of multiple functional units,
4. *suspense*: the delay between the adoption of an intention and its dropping can be measured by either the number of actions executed, or the number of reasoning cycles completed, in the mean time,
5. *dynamic points*: a violation of an agent's expectations occurs each time a plan-failure or a false belief is detected.

This reformulation of the principles shows that they can be captured computationally using agent internal events regarding actions, emotional valencies, goals, plans, beliefs and functional units. Of these concepts only the latter has not yet been addressed, which will be remedied in the following section.

Plot Units and Graphs

Functional plot units were introduced by Lehnert (1981) in the context of plot summarisation and rely on a directed,

[3] Berov (2017) represents emotions using vectors in a space that contains a valence dimension. Details on the mapping can be found in (Gebhard 2005; Alfonso, Vivancos, and Botti 2017).

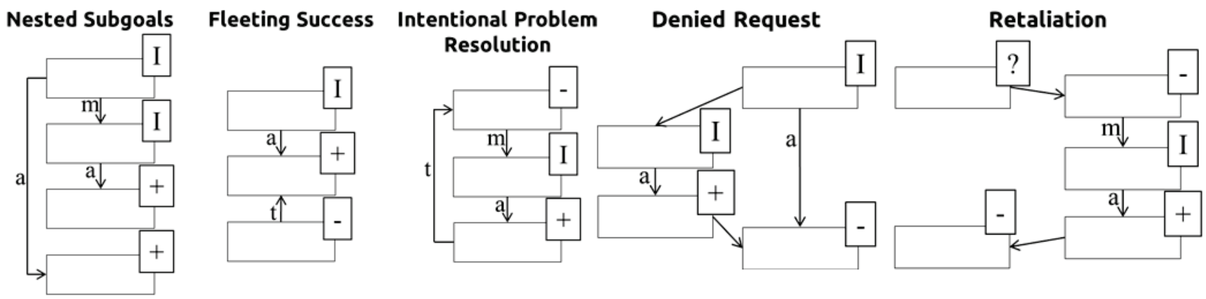

Figure 1: Examples of complex plot units; '?' is a wildcard for arbitrary vertex type. Adopted from (Lehnert 1981).

chronologically ordered graph representation of plot, which will be introduced here along the way. The need for such a representation is not problematic, because in the present use case no analysis of natural language stories is required to obtain it. Rather, the graph can be generated automatically from internal and external events in the simulation.

In Lehnert's graph formalism vertices are used to denote events, while edges represent their logical connection. For each character a subgraph is created that contains their respective, subjective evaluation of events. These subgraphs are connected in ways specified below, if an event affects several characters. The following description can be best read in conjunction with fig. 2 for a practical example of the formalism in use.

Each subgraph can contain three types of vertices:

- **+**: an internal or external event that is appraised with a positive emotion by the character,
- **–**: an internal or external event that is appraised with a negative emotion by the character,
- **I**: a neutral internal event that denotes the setting of an intention by the character[4].

These vertices can be connected by five types of edges:

- **m**: a motivation edge can lead from any type of vertex to an `I` vertex and denotes that the former event causes the intention,
- **a**: an actualization edge can lead from an `I` vertex to a + or − vertex and denotes that the intention causes the latter event,
- **t**: a termination edge can connect either two non-neutral, or two `I`, vertices and leads from the event that terminates an affective/intensional state to the one that is terminated (respectively supplanted in the `I` case),
- **e**: an equivalence edge can connect two non-neutral, or two `I` vertices, and denotes that an event has multiple affective evaluations, or two intentions are congruent; the edge direction is anti-temporal,
- a cross-character edge is temporally directed and can connect all types of vertices, so long as they belong to different characters' subgraphs; it denotes that an event is affecting several characters, prominently also including speech acts.

Using this notation, Lehnert identifies complex plot units that "represent general plot configurations" (Lehnert 1981). It is these units that are referred to in the last section, when a representation of the principle of functional poly-valence is suggested. Lehnert introduces nine complex intra-character units, and over 20 complex cross-character units, a sample of which is given as example in fig. 1. However, there is no theoretic limit to the size of complex units, so more functions can be derived e.g. from existing narratives. To illustrate the formalism in action, an abridged manual analysis of the plot of a fable is provided in fig. 2.

All the information necessary to automatically generate these plot-graphs is present in the suggested character architecture. When an intention is selected by an agent, an `I` vertex is added to its subgraph. For other internal events, the affective appraisal is consulted first. If an event causes an emotion it is included in the character subgraph and the vertex's type can be inferred from the emotion's valence. Edge-types can be derived based on the following considerations:

- **m**: is added when processing an event causes the agent to select a desire as an intention,
- **a**: is added to connect an intention with each action that is executed by the agent while it is on the top of its intention stack,
- **t**: is added either when processing an event causes a change in the agent's belief base, or when executing a plan to fulfill an intention explicitly involves dropping another intention,
- **e**: is added only when appraising an event results in several emotions, in which case an appropriate vertex is added for each emotion; there is no need for equivalence edges between `I` vertices, if each vertex is simply allowed to be the source and target of several edges.

Once a simulation run is concluded, detecting complex plot units in the graph is an instance of the subgraph isomorphism problem, for which solutions exist (Ullmann 2011). This allows identifying events that are part of several plot units in order to detect functional poly-valence. It is inter-

[4] In the original work, this type of node is called `M` for mental event. However, since they are only ever used to represent character goals it seems helpful to translate them into the terminology used here.

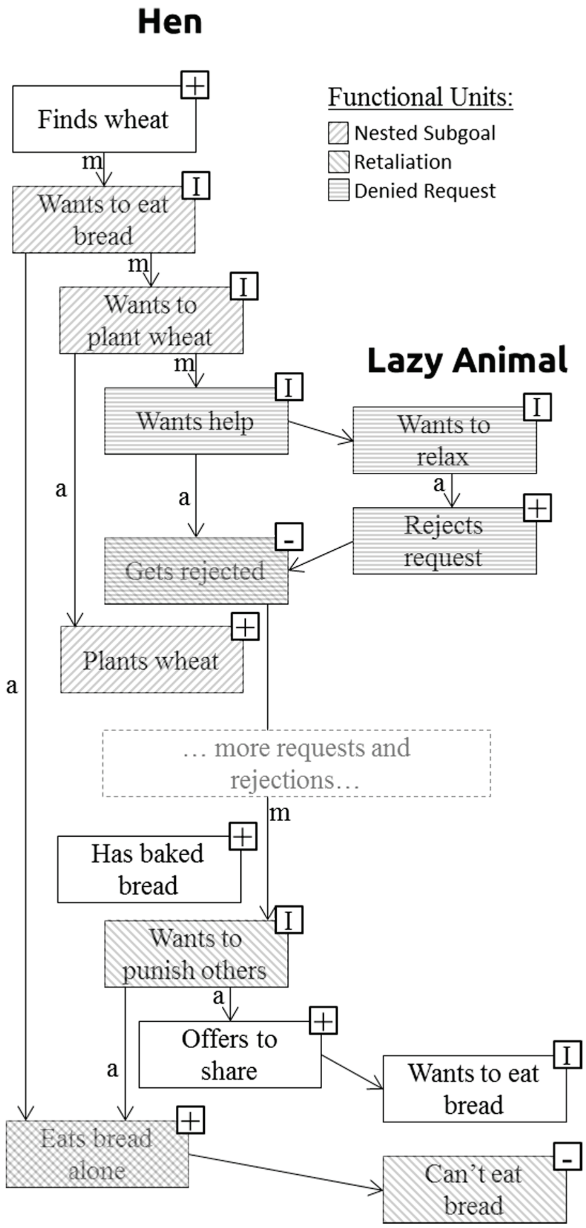

Figure 2: Abridged, manual analysis of the plot of the fable "The Little Red Hen": In the fable a hen and three lazy animals live on a farm. The hen finds a grain of wheat and decides to make bread. She repeatedly asks the other animals for help but always gets rejected. In the end she makes the bread all by herself, asks the other animals if they want to help her eat it, but when they accept she eats the bread alone. The three animals and the repetitions of the denied requests are compressed for brevity. Note the functional poly-valence of 'gets rejected' and 'eats bread alone', which participate in multiple plot units.

esting to note, that Lehnert's representational formalism explicitly models not (only) actions but mostly internal events like intentions and emotions. This means that, apart from tellability analysis, it also offers itself as a general-purpose plot-graph for systems operating on an embedded narrative aware notion of plot (at least in its Palmerian sense).

Ryan (1991, p. 218ff) perspicaciously discusses several drawbacks of Lehnert's approach and puts forward her own formalism that is better capable of capturing certain narrative phenomena, especially related to representing embedded narratives. However, her suggestion does not seem to yield itself to computational modelling as readily as Lehnert's.

Related Work

The approach outlined in the present paper is in the vein of (Pizzi and Cavazza 2007) in that it also aims at relating "the character's psychology to aesthetic elements of the narrative". However, instead of taking a narrative approach like Pizzi and Cavazza it explores the cognitive route to modelling affect. The decision to do so rests on the work of narratologist Alan Palmer (2004), who describes fictional minds in terms of 'real minds' theory of mind. Both approaches are alike in that they attempt to move away from an action/plan-only notion of plot. However, they also differ in that the suggestion outlined here relies rather on virtuality as the subjective experience of actions, while Pizzi and Cavazza's work puts more emphasis on a character's emotional progression. It is interesting to observe, how such differences in narratological framework relate to differences, but also similarities, in computational modelling. For instance, Pizzi and Cavazza search for patterns in the emotional arc ("state heuristic") of a character, while here a search for subgraphs in a plot-graph is proposed. Yet the pattern labeled "shattered hopes" by Pizzi and Cavazza seems to encode a structure quite similar to the complex plot unit "fleeting success" introduced here (see fig. 1).

Another interesting point of comparison for the approach outlined here is the Glaive system (Ware and Young 2014). On the surface the differences are striking: Glaive employs a central state-space planner to compute a plot-as-solution, while the present work relies on a multi-agent simulation system for an emergent plot. Grounded in their comparable narratological ancestry the systems, however, share a fundamental understanding of the nature plot. Glaive's plot-representation includes not only the solution for the search-problem but also all paths branching from it in the state-space that represent goal-directed character behavior. As Ware and Young point out, their notion of plot thus contains failed plans, which under the label of prospective embedded narratives are also included in the plot-understanding proposed here. While the present paper also includes retrospective embedded narratives—mainly affect—into its plot representation, it is important to emphasize the shared focus on the virtual in addition to the actual. This virtuality is instrumented for different purposes by the two approaches. Glaive uses it to ensure intentionality-based consistency, while the present paper suggests to use it to estimate plot aesthetics.

The notion of embedded narratives is implicitly present in other systems as well. Chang and Soo (2008) model social influence in narratives by implementing a system that allows agents to reason about other agent's reasoning. Thus, agent's plans can happen to include the prospective plans of other agents—a classic case of embedded prospective narratives. Discussing an alternative to their original modelling of Othello, Chang and Soo observe: "with the new rule above, Iago can easily ask Othello to kill Desdemona and complete the task of the original story, losing most of the intrigue". This sentiment is understandable using the reasoning outlined above in the second section: the new plot includes only one embedded narrative while the original plot required three, namely one for Emilia, one for Cassio and one for Othello. The network of the virtual is more complicated in the latter case, which is the precise reason why it appears more tellable.

Also Brenner (2010) includes embedded narratives into plot: "Since during a continual planning episode usually multiple plans are being generated, executed, and revised, we consider as the plot the execution history of the episode, annotated with relevant (possibly false) beliefs and goals".

Discussion

The present paper addresses the problem of determining the quality of a plot represented in a computational narrative system. For that it introduces Ryan's (1991) notion of tellability as a function of embedded narratives: prospective and retrospective virtual representations of narrative events by affected characters. This relation is condensed by Ryan into a pre- and a post-condition of well formed plot, as well as five formal principles correlating with higher tellability and independent of content or context, which makes them good potential guidelines for computational storytelling systems.

Based on this narrative theory, the paper proposes a computational model in the context of an affective BDI simulation. Under such an approach, embedded narratives are represented using the internal events from the runtime execution of character agents' reasoning cycles. This, in turn, allows reformulating Ryan's conditions and principles in the terms of the character architecture and by that make them accessible for a computational system.

One of the tellability principles requires a plot-functional analysis of narrative events. For this purpose the paper suggests using Lehnert's (1981) complex plot units, which rely on analysing structures in a specific plot-graph representation introduced by her. It is argued that the introduced character architecture is sufficient to generate Lehnert's plot-graphs, and that known solutions to the subgraph isomorphism problem can be employed to perform the analysis required by the tellability principle. The other principles can be represented using internal events of the agent's reasoning cycles.

However, an implementation of the outlined computational model is far from trivial. For instance, an environment-model is required that is both procedural—in that it needs to interact with individual agents—as well as declarative, in that epistemic conflicts in agent's beliefs need to be identified. This, as well as inter and intra-agent conflict detection also requires a reasoning system apart from individual agent's reasoning cycles but with access to their internal data. Furthermore, to implement a quantifiable tellability measure a mathematical formalisation of the intuitions captured by the tellability principles is required. Such a formalisation needs to also address the issues of aesthetic balance that are only cursory discussed in the background literature.

Thus, the purpose of this publication is to introduce the theoretical model in order to be able to invite commentary, before continuing with its realisation. It aims at starting a discussion of the outlined narrative framework, especially on the role of embedded narratives in plot and with regard to already existing narrative systems. It also invites collaboration on possible implementations and formalisations of the framework, again especially in combination with already existing approaches.

While the ideas reported here concentrated on plot, connections can also be made to the automatic generation of discourse. For instance, if a tellability-analysis determines that most of the plot's quality stems from the subgraph of a single character, a focalization from this character's point of view can be considered (see e.g. Fludernik 2009, p. 153). This would require including only those events from the plot-graph that are connected to the focalizer's subgraph. The generation of internal events discussed above would allow to generate the discursive access to a protagonist's consciousness typical for internal focalization (Fludernik 2009, p. 159). Focalization seems thus to be a promising avenue for further investigation.

Acknowledgements.

The author is grateful for support for this work provided by an Alexander von Humboldt Ph.D. fellowship funded by an Anneliese Maier-Forschungspreis awarded to Mark Turner.

References

Abbott, P. 2014. Narrativity. In Hühn, P.; Meister, J. C.; Pier, J.; and Schmid, W., eds., *Handbook of Narratology*. Berlin: De Gruyter.

Alfonso, B.; Vivancos, E.; and Botti, V. 2017. Toward formal modeling of affective agents in a BDI architecture. *ACM Transactions on Internet Technology (TOIT)* 17(1):5.

Berov, L. 2017. Steering plot through personality and affect: an extended BDI model of fictional characters. In *KI 2017: Advances in Artificial Intelligence: Proceedings of the 40th Annual German Conference on AI*. Dortmund, Germany: Springer International Publishing. Forthcoming (accepted).

Brenner, M. 2010. Creating dynamic story plots with continual multiagent planning. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*}, AAAI'10, 1517–1522. Atlanta, Georgia: AAAI Press.

Bruner, J. 1991. The narrative construction of reality. *Critical inquiry* 18(1):1–21.

Chang, H.-M., and Soo, V.-W. 2008. Simulation-based story generation with a theory of mind. In *Proceedings of the Fourth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 16–21.

Fludernik, M. 2009. *An Introduction to Narratology*. Routledge.

Gabriel, R. P. 2016. in the control room of the banquet. In *Proceedings of the 2016 ACM International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software*, 250–268. ACM.

Gebhard, P. 2005. ALMA: a layered model of affect. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, 29–36. ACM.

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

Labov, W. 1972. *Language in the Inner City: Studies in the Black English Vernacular*, volume 3. University of Pennsylvania Press.

Lehnert, W. G. 1981. Plot units and narrative summarization. *Cognitive Science* 5(4):293–331.

Palmer, A. 2004. *Fictional Minds*. U of Nebraska Press.

Pizzi, D., and Cavazza, M. 2007. Affective storytelling based on characters' feelings. In *Intelligent Narrative Technologies: Papers from the AAAI Fall Symposium*, 111–118.

Prince, G. 2003. *A Dictionary of Narratology*. U of Nebraska Press.

Rao, A. S., and Georgeff, M. P. 1995. BDI agents: from theory to practice. In *Proceedings of the 1st International Conference of Multiagent Systems*, 312–319.

Ryan, M.-L. 1991. *Possible Worlds, Artificial Intelligence, and Narrative Theory*. Indiana University Press.

Ryan, M.-L. 2013. Possible worlds. In Hühn, P.; Meister, J. C.; Pier, J.; and Schmid, W., eds., *the living handbook of narratology*. Hamburg: Hamburg University.

Ullmann, J. R. 2011. Bit-vector Algorithms for Binary Constraint Satisfaction and Subgraph Isomorphism. *Journal of Experimental Algorithmics* 15.

Ware, S. G., and Young, R. M. 2014. Glaive: A State-Space Narrative Planner Supporting Intentionality and Conflict. In *Proceedings of the Tenth Artificial Intelligence and Interactive Digital Entertainment Conference*.

Wei, H. 2010. Embedded narrative in game design. In *Proceedings of the International Academic Conference on the Future of Game Design and Technology*, Futureplay '10, 247–250. Vancouver, British Columbia, Canada: ACM.