Intelligent Narrative Technologies and Social Believability in Games:
Papers from the AIIDE 2015 Joint Workshop

A Tripartite Plan-Based Model of Narrative for Narrative Discourse Generation

Camille Barot and **Colin M. Potts** and **R. Michael Young**
Liquid Narrative Group, Department of Computer Science
North Carolina State University, Raleigh, NC 27695
cbarot@ncsu.edu, cmpotts@ncsu.edu, young@csc.ncsu.edu

Abstract

This paper presents a tripartite model of narrative in which three different kinds of actions are modeled: story actions, discourse actions, and narration actions. We argue that this model can support particularities of narrative discourse where typical discourse models do not (e.g., ambiguity/change of beliefs), and accounts for the difference in expressivity of the media used to convey the narrative, as opposed to bipartite models.

Introduction

Narratologists typically distinguish two aspects of narratives: story and discourse (Chatman 1980). The story is a conceptualization of the world of the narrative, with the characters, actions and events that it contains, while the discourse is composed of the communicative elements that participate in its *telling*. Research on computational models of narrative has produced many models of story, based for instance on plans (Riedl and Young 2010), ontologies (Swartjes and Theune 2006), or conceptual graphs (Elson and Mckeown 2007), but little attention has been paid to building integrated models of narrative discourse. In contrast, research in computational linguistics has led to a variety of formal models of discourse, used to represent discourse such as explanatory texts (Maybury 1991), task-oriented dialogues (Lambert and Carberry 1991), or advisory dialogues (Moore and Paris 1993). We argue that these models are not expressive enough to support the representation of narrative discourse.

Moreover, we argue that the bipartite model of narrative is too limited when considering narrative across a range of media. We believe that a differentiated representation of authorial intentions in terms of discourse acts and of realization choices for these discourse acts in different media would allow automated systems to account for specific medium constraints in the generation of a narrative, e.g., text versus movies.

This paper presents a tripartite model of narrative that distinguishes between a medium-agnostic discourse layer and a medium-specific narration layer. This model is used as a basis for the representation of story, discourse and narration knowledge in a narrative generation framework.

Narrative Discourse

While the bipartite model of narrative (Chatman 1980) has been advanced in the field of narrative generation (e.g., (Young 2007)), several narratologists have argued for representing narrative as a tripartite construct. According to Genette (1980), a narrative is composed of three distinct layers: the story, comprising the actions and states of the characters and the world in the chronological order in which they happen; the discourse, which contains an organization of the story elements in the order in which they are presented to an audience; and the narration, which is the act of narrating taken in itself, in which the discourse is conveyed through a particular medium. However, the discourse layer is not simply an ordered subset of elements of the story layer. Genette argues that every discourse implies a narrator. In this, the discourse is an intentional structure through which the narrator "regulates the narrative information" given to the audience, and its representation should include these intentions.

Some of the discourse models used in computational linguistics acknowledge intentional, attentional and rhetorical structure (e.g., (Moore and Paris 1993), (Grosz and Sidner 1986)), but they do not capture those dynamics of beliefs that are specific to narrative discourse. Pratt (1977) argues that narratives, as opposed to other types of discourse, fall under a particular discourse situation, where the hearer takes the role of an *Audience*, and agrees on a tacit contract with the speaker to give up her speech turn in exchange for certain quality expectations regarding the narrative. This means that the goal of the author of a narrative discourse is not only to convey a sequence of events, but also to entertain the audience and convince it of the interest of the story throughout the telling, so that the contract is maintained. In our view, this goal manifests itself through the deliberate fostering of narrative effects such as suspense or surprise that stem from the careful preservation of ambiguity in the discourse. In this sense, the narrative discourse is carefully designed for the audience to revise their mental model of the story as they progress through the narrative. Consider for instance the following text:

> *There are three characters named Blue, White and Red. Red and White are at the manor, and they get into a fight. Later on, White gets killed. Red wanted to avoid White, so he had left the manor prior to the murder. White had been blackmailing Blue, because he wanted*

to get rich. This made Blue angry. Blue wanted White dead, so he arrived at the manor and he killed him.

This narrative is designed in such a way that after receiving the utterance *"White gets killed,"* the audience would hold false beliefs about the story, inferring that Red is the murderer. However, after the utterance *"Red had left the manor,"* the audience is expected to drop this belief.

Existing discourse models, whether based on Rheotorical Structure Theory (Mann and Thompson 1988) (e.g., (Hovy 1990), (Moore and Paris 1993)) or on AI planning (e.g., (Appelt 1985), (Young and Moore 1994)) assume a monotonic progression of the audience's beliefs about the discourse's propositional content. To account for narrative discourse, there is a need for a model that would consider the possible negation of these beliefs. This is particularly important in cases where certain beliefs shouldn't hold for a narrative effect to be achieved: in the previous example, learning that Red has left the manor and thus cannot be the murderer only comes as a surprise if the audience doesn't already suspect that Blue is the murderer. To generate such a narrative, a system would have to reason not only on the inferences that should be made by the audience to create the desired beliefs, but also on the potential for inferences that would thwart the narrative effects, so that those can be anticipated and avoided. To do so, an explicit representation of an audience's belief changes and inferences is needed.

Genette's tripartite model also calls for a separation of the medium-independent discourse structure from the medium-dependent narration choices. Most narrative generation approaches either do not make the distinction (e.g. (Jhala and Young 2010)), or create the two layers in a pipeline approach (e.g., (Bae and Young 2014)). Yet, the selection and organization of story content to be included in a discourse cannot be done without considering the medium in which the narrative will be realized, as the specificities of the medium can make certain story elements more difficult or even impossible to convey. For instance, the common trope where two different characters turn out to be the same person[1] can be easily realized in text, but requires some adaptation to be done in movie form as the audience would visually recognize the actor playing both. It is thus necessary to consider medium-independent and medium-specific knowledge together in an integrated process, rather than in a pipeline approach. However, by segregating medium-independent elements of knowledge from their medium-specific counterparts, one allows for the reuse of narrative discourse knowledge across different media. For instance, in a cinematic discourse generator, a strategy for conveying the emotion of a character could be encoded as a decomposition rule $\texttt{show-sadness(}c\texttt{)} \rightarrow \texttt{close-up-shot(}c,cry\texttt{)}$. When adapting the discourse generator to create textual narratives, the rule would have to be entirely rewritten for the new medium: $\texttt{state-sadness(}c\texttt{)} \rightarrow \texttt{text-clause(}c,cry\texttt{)}$. The separation of discourse and narration knowledge would allow for such a rule to be encoded in a medium-independent way (e.g., $\texttt{convey(}sad(c)\texttt{)} \rightarrow \texttt{convey(}cry(c)\texttt{)}$) and reused for generating narrative in various media. Discourse-level rules would capture rhetorical relations, as opposed to medium-specific idioms.

[1] "Two Aliases One Character" - http://tvtropes.org/

Limitations of Previous Approaches

Early work on content planning for Natural Language Generation (NLG) followed two main approaches. Work grounded in AI planning (e.g., (Cohen and Perrault 1979), (Appelt 1985)) formalized speech acts as plan operators having preconditions and effects on the hearer's beliefs. Work by Hovy (1990) and Maybury (1992) operationalized schemata from Rhetorical Structure Theory as decompositional patterns. Later work by Moore and Paris (1993) and Young and Moore (1994) combined a representation of preconditions and effects of discourse actions with rhetorical patterns in order to preserve in the construction of the discourse an explicit representation of both the intentional and rhetorical knowledge. The latter uses a representation based on Decompositional Partial Order Causal Link (DPOCL) plans (Young, Pollack, and Moore 1994), which allows the algorithm that generates the discourse to plan for desired audience inferences through their encoding in decomposition schemata. However it does not account for potential undesired inferences that might threaten causal links, i.e. negate conditions needed for further discourse actions.

Other researchers have tackled the problem of narrative-specific discourse generation. Some of their systems are medium-agnostic and model discourse as a reordered and filtered subset of the events in a story plan (e.g., (Niehaus and Young 2014), (Cheong and Young 2014), (Bae and Young 2014), (Hoek, Theune, and Linssen 2014)). These systems ignore part of the problem of content selection, as they assume an implicit one-to-one mapping between a partial story plan and the resulting narrative. Others (e.g., (Theune, Slabbers, and Hielkema 2007), (Rishes et al. 2013)) have used explicit or learned mapping rules from formal story representations to lexico-syntactic representations, but they suffer from the same limitation. Similarly, the Author system (Callaway and Lester 2002) can create rich narrative text from a representation of story content as a set of "narrative primitives," but does not link those to a representation of the story itself. The former narrative discourse generators could be used with the latter narrative text realizers in a pipeline model, however this approach would not allow the medium constraints to be taken into account during the selection and organization of the discourse content when addressing media other than text. The Curveship narrative text generation system (Montfort 2011) deals with both content selection and text realization to generate discourse variations from the same story according to parameters for focalization, temporal ordering, or pace. However the variations come directly from settings entered by a user and no reasoning is done on their effects on audience's beliefs on the story.

The Darshak system (Jhala and Young 2010) extends DPOCL to plan sequences of camera shots for automatically generated animated movies. It uses a representation of cinematic idioms — commonly-used sequences of shots used to film story events, such as *Show Conversation* or *Show Gunfight* — as decomposition schemata. However Darshak does not propose a distinction between discourse and narration elements, as it addresses cinematic discourse specifically, and its representation of the audience's mental model is too limited to support reasoning on desired or undesired inferences.

Tripartite Model of Narrative

Following Genette's tripartite model, our representation of a narrative is separated into three layers: story, discourse and narration. Each layer is composed of actions that form causal chains leading to the realization of story goals, communicative goals, and narration goals, respectively. Actions from the story layer correspond to actions executed by story world characters. The structure of characters goals and intentions is expressed through groupings of these actions as *frames of commitment*. Actions from the discourse layer are communicative actions whose effects represent the speaker's intention to affect the mental state of the audience. These actions can be considered as speech acts (i.e., they can be directly mapped to utterances at the narration layer) or can correspond to rhetorical goals and be decomposed further at the discourse layer into sub-actions. Actions from the narration layer correspond to the realization of surface speech acts, and have surface effects that characterize the fact that some particular content has been successfully expressed in the medium (e.g., *told*, *shown*). These actions can be grouped into medium-specific patterns such as cinematic idioms.

The actions of the discourse and narration layers are grouped in decomposition frames, each frame corresponding to the realization of a more abstract communicative action. However, the same discourse or narration action can belong to several of these frames if it participates in achieving several discourse goals. For instance, a text clause describing an action can fulfill a discourse goal of informing the audience about the existence of the character performing the action, of informing the audience of a character property that is displayed by the action, or of informing the audience of the occurrence of the action itself. Similarly, the propositional content of the discourse actions corresponds to elements of the story layer, but several discourse actions can refer to the same story element.

An example of a narrative model corresponding to an extract of the text previously introduced is shown in Figure 1.

Knowledge Representation

We describe here our knowledge representation for the elements of the tripartite model of narrative.

Story Layer

The story layer representation is based on IPOCL plans (Riedl and Young 2010): partially-ordered, causally-linked sets of steps corresponding to character actions. This representation extends that of classical planning by including explicit representations of character motivations and plans. Each story takes place in a specific story domain:

Definition 1 (Story Domain) *A story domain Δ is a set of story action schemata of the form $(\alpha, V, A, Pre, \mathit{Eff})$ where α is a unique action identifier, V is a set of variable identifiers used as parameters of the action schema, $A \subset V$ the subset of these variables corresponding to the actors who have to consent to the action, and Pre and Eff are sets of literals corresponding respectively to preconditions and effects of the action schema.*

A story problem corresponds to a specific "setting" in a story world. It defines the characters and objects that are instantiated in a particular subset of stories for a given domain, as well as the initial state of the story world. A story problem also defines a "goal" state, which corresponds to a set of facts that have to be true at the end of the story.

Definition 2 (Story Problem) *A story problem Π is a tuple (Δ, C, I, G), where Δ is a story domain, C is a set of constants, with $\Lambda \subset C$ the subset of constants corresponding to actors, I is a set of literals that are true in the initial state and G is a set of literals that corresponds to the goal state.*

Definition 3 (Step) *A step s is a tuple (σ, α, B) where σ is a unique step identifier, α is an action identifier of an action schema in Δ and B is a set of parameter bindings of the form (v, c) where $v \in V$ and $c \in C$.*

A story plan is thus defined as the following:

Definition 4 (Story Plan) *A story plan P is a tuple (S, O, L, F), where S is a set of steps, O is a set of temporal orderings of the form $(s_1 < s_2)$ where $s_1, s_2 \in S$, L is a set of causal links of the form (s_1, l, s_2) where $s_1, s_2 \in S$, and l is an effect of s_1 and a precondition of s_2, and F is a set of frames of commitment.*

Frames of commitment correspond to the interval of intentionality during which a character tries to achieve a goal through a particular subplan:

Definition 5 (Frame of Commitment) *An frame of commitment f in a plan P is a tuple $(\lambda_f, g, m, S_f, s_f)$ where λ_f is an actor, g is a literal corresponding to the goal of the frame of commitment, $m \in S$ is a motivating step such that m has the effect $\texttt{(intends}\ \lambda_f\ \texttt{g)}$, $S_f \subset S$ is a subplan such that $\forall s_i = (\sigma_i, \alpha_i, B_i) \in S_f$ with $(\alpha_i, V_i, A_i, Pre_i, \mathit{Eff}_i) \in \Delta, \exists a_i \in A_i$ such that $(a_i, \lambda_f) \in B_i$ (in other words, the actor of the frame must consent to all the steps in S_f), and $s_f \in S_f$ is the final step of the frame of commitment such that $g \in \mathit{Eff}_{s_f}$.*

Discourse Layer

The discourse layer representation is based on a DPOCL representation (Young and Moore 1994), extended to represent additional aspects of the audience's mental state and inferences. Each step of a discourse plan correspond to a communicative action whose propositional content is a part of a story plan: In order to be referenced in the constraints and propositional content of the discourse operators, the story plan representation is mapped to a set of plan description formulae, following the approaches used by Ferguson (1992) and Jhala and Young (2010). The reified story plan is defined as follows:

Definition 6 (Reified Story Plan) *A reified story plan is a tuple (C_D, Φ_D) with C_D a set of constants denoting reified story elements and Φ_D a set of story description formulae.*

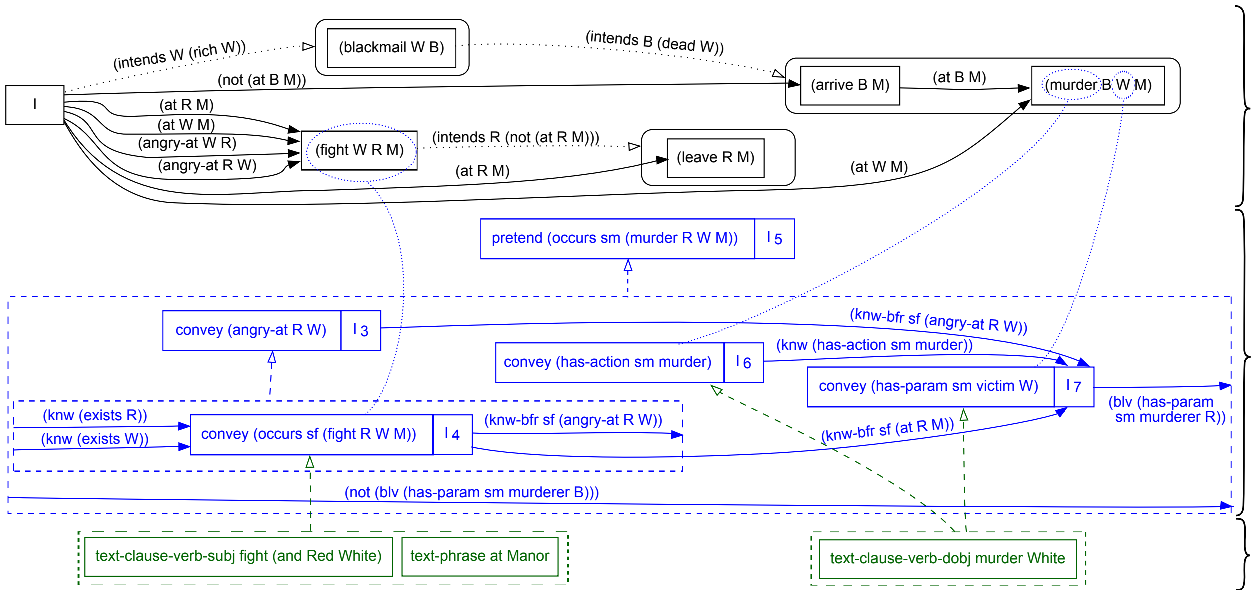

Figure 1: Example of Tripartite Model of Narrative. This graph shows the three layers: story (black, top), discourse (blue, middle), and narration (green, bottom). Actions, frames of commitment, and decomposition frames are denoted by solid rectangles, rounded rectangles, and dotted rectangles respectively. Causal links are solid arrows labeled by their proposition. Intention links are dotted lines and decomposition links are dotted arrows. Dotted blue lines denote reifications of story elements. Discourse actions have an additional label, I_k, that names the inference frame associated with the action. Causal links can come to or from either the main body (the action's effects) or the inference frame (the audiences' inferences).

The set of constants C_D for a given story plan P and problem Π contains reified version of step identifiers σ, characters $\lambda \in \Lambda$, non-character constants $c \in C \setminus \Lambda$, frames of commitment $f \in F$, action identifiers α and variable identifiers v. The set of story description formulae Φ_D results from a mapping from story plan P and problem Π to the following predicates, with l representing a literal:

(has-action σ α)
(has-parameter σ v c)
(has-actor σ λ)
(has-precondition σ l)
(has-effect σ l)

The causal and temporal structure of the story plans is represented by two predicates:

(causal-link σ_1 l σ_2)
(comes-before σ_1 σ_2)

Finally, the intentional structure of character plans is represented by the following predicates:

(has-goal f l)
(has-intention λ f)
(motivates σ f)
(is-in-plan σ f)

Audience Mental State Discourse actions have effects on the expected mental state of the audience. This mental state is represented by a set of formulae that are either *known* or *believed* to hold. For now our model assumes a reliable narrator, so the story elements conveyed by the discourse actions are considered known. Knowledge formulae can be added to the mental state, but never negated, and must be consistent with the story layer. Beliefs are inferred by the audience, and can be inconsistent with the story layer (but not with knowledge or other beliefs). They may be subsequently removed based on new beliefs or knowledge. The evolution of the mental state of the audience along with the discourse is represented by a set of *inference frames*, one per discourse action. A mental state also contains inference rules, which inform the inference process. These are if-then pairs of formulae and represent possible inferences.

Definition 7 (Inference Frame) *An inference frame i in a discourse D is a tuple (K, B) where K is a set of knowledge formulae and B is a set of belief formulae.*

Predicates used in belief and knowledge formulae can be either: *existence* predicates on story characters λ or constants c, *occurrence* predicates on reified story steps σ or frames of commitment f, story description predicates, or *state* predicates used in the story layer. State predicates can be indexed according to a simple temporal model: A formula can be believed/known to hold immediately before or at a story step σ. From each inference frame, a partial story plan can be reconstructed (c.f. Figure 2).

Discourse Actions The actions of the discourse layer are instantiated from schemata defined in a discourse domain.

Definition 8 (Discourse Domain) *A discourse domain Δ_D is a set of discourse action schemata of the form $(\alpha_D, V_D, K_D, Inv_D, Pre_D, Eff_D)$ where $\alpha_D = (a, p)$ is a unique action identifier composed of an action type a and a proposition type p , V_D is a set of variable identifiers used*

```
(know (occurs sf (fight Red White Manor)))
(know-before sf (angry-at Red White))
(know-before sf (angry-at White Red))
(know-before sf (at Red Manor))
(know-before sf (at White Manor))
(know (has-action stepm murder))
(know (comes-before stepf stepm))
(know (occurs stepm
      (murder ?murderer White ?location)))
(know-before sm (at ?murderer ?location))
(know-before sm (at White ?location))
(believe (= ?location Manor))
(believe-before sm (at Red Manor))
(believe-before sm (angry-at Red White))
(believe (= ?murder Red))
```

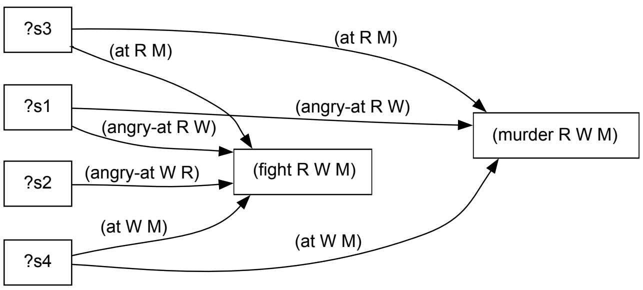

Figure 2: Inference Frame I_7: knowledge and beliefs of the audience, and graph of the inferred story.

as parameters of the action schema, K_D is a set of constraints on the application of the action schema, Inv_D is a set of invariants that have to hold during the execution of the action and Pre_D and Eff_D are sets of non-ground literals corresponding respectively to preconditions and effects of the action schema.

We consider two action types depending on the truth value of the proposition that is asserted: `convey` or `pretend`. Proposition types correspond to the different natures of predicates: existence, occurrence, plan description element or state. The variables used as parameters by the schema can be identified with either reified story elements, formulae with story predicates, or story description formulae, depending of the constraints of the schema. Constraints are conditions that have to hold in the story layer for the action schema to be applicable, while preconditions have to hold on the mental state of the audience. Examples are given in Table 1.

These actions are linked to sub-actions in the discourse or narration layer through decomposition frames that instantiate decomposition schemata. These schemata correspond to a particular strategy to realize a given discourse action.

Definition 9 (Discourse Decomposition Schema) *A discourse decomposition schema is a tuple $(\alpha_D, V'_D, K'_D, Inv'_D, S_D, O_D, L_D)$ where α_D is an action identifier, V'_D is a set of variable identifiers used as parameters of the decomposition schema, such that $V'_D \subset V_D$ for the corresponding action schema, K'_D is a set of constraints on the application of the decomposition schema, Inv'_D a set of invariants, S_D is a set of tuples (α_i, B_i) corresponding to actions schemata and binding constraints on the variables of these actions schemata, forming the steps of the decomposition's subplan, and O_D and L_D are respectively temporal ordering and causal links on the steps of this subplan.*

Act:	`convey`
Prop.:	`comes-before`
Param.:	$\sigma_1, \sigma_2 \in C_D$
Const.:	(comes-before σ_1 σ_2)
Prec.:	(knows (occurs σ_1)) (knows (occurs σ_2))
Effects:	(knows (comes-before σ_1 σ_2))

Act:	`pretend`
Prop.:	`has-parameter`
Param.:	$\sigma, v, c \in C_D$
Const.:	(not (has-parameter σ v c)) (has-parameter σ v c')
Inv.:	(not (believes (has-parameter σ v c')))
Effects:	(believes (has-parameter σ v c)))

Table 1: Examples of Discourse Action Schemata

An example of decomposition schema is given in Table 2.

Decomp.:	*Convey precondition of action*
Act:	`convey`
Prop.:	*state*
Param.:	l, σ
Const.:	(has-precondition σ l) (occurs σ)
Steps:	(convey (occurs σ))
Effects:	(knows-before σ l)

Table 2: Example of Discourse Decomposition Schema

The discourse steps and plans are thus defined as follows:

Definition 10 (Discourse Step) *A discourse step is a tuple $(\sigma_D, \alpha_D, B_D, I_{\sigma_D})$ where σ_D is a step identifier, α_D is an action identifier of a discourse action schema in Δ_D, B_D is a set of parameter bindings from the variables in V_D to elements of $C_D \cup \Phi_D$ or to formulae with story predicates, and I_{σ_D} is the associated inference frame.*

Definition 11 (Discourse Plan) *A discourse plan P_D for a story plan P and a discourse domain Δ_D is a tuple (S_D, O_D, L_D, F_D), where S_D is a set of discourse steps with unique identifiers, O_D is a set of temporal orderings on the steps in S_D, L is a set of causal links between the steps in S_D, and F_D is a set of decomposition frames $f = (\sigma_D, \delta, S'_D, B'_D)$ where σ_D is the identifier of the abstract step that subsumes the frame, δ is a decomposition schema, S'_D is a set of steps that belong to the frame, and B'_D a set of binding constraints on the variables of the decomposition schema.*

Narration Layer

The definitions of narration domains, decomposition schemata, steps, and plans mirror those of their discourse-level counterparts. The action types, however, are not speech acts but medium-specific items, such as `text-clause` or `camera-shot`, that can then be transformed into the final narrative by a medium realizer. Moreover, the decomposition frames of the narration layer are tuples $(\sigma_D, \delta_N, S'_N, B'_N)$ where σ_D is an identifier of a step in the discourse plan, and S'_N is a subset of the steps of the narration plan: The decomposition schemata of the narration domain correspond to medium-specific idioms and decompose abstract discourse actions into concrete narration actions. For instance, the discourse action of conveying that a character has a dominant personality could be decomposed into a camera shot of the actor from a low angle.

Narrative Generation

This tripartite model serves as a basis for a planning-based narrative generation process that generates a narrative from an existing story representation. This algorithm is based on the DPOCL planning algorithm (Young, Pollack, and Moore 1994). It takes as input a specification of a planning domain as a set of discourse and narration action and decomposition schemata, a specification of a planning problem containing the reified story plan as an initial state, a set of inference rules for the audience, and a set of "islands" that serve as discourse goals (Riedl et al. 2008). These islands are partially ordered sets of beliefs that the author would want the audience to hold while they experience the narrative. For instance, the narrative presented in this paper could be generated from achieving successively `(believe (has-parameter stepm murderer Red))` and `(know (has-parameter stepm murderer Blue))`. To create a narrative plan, the planner adds new discourse and narration steps to the plan, using the patterns specified in the decomposition schemata, until all preconditions are satisfied and all discourse actions have been decomposed to the level of narration actions. The difference from the classical DPOCL algorithm is that, for each step being added, the planner will recompute the inference frame for each step in the plan, by applying the inference rules based on possible linearizations of the steps up to it. This allows the planner to detect potential threats to causal links between steps when the opposite beliefs are created or negated by the inferences.

These rules can be both generic and domain specific. A more specific rule would be

```
((if (and
       (blv (occurs ?s (murder ?k ?v ?l)))
       (blv-bfr ?s (at ?p ?l))
       (blv-bfr ?s (angry-at ?p ?v))))
 (then (blv (= ?k ?p))))
```

which is used in the example presented in Figure 1 to simulate the audience inferring that *Red* is the killer. This rule also indicates that the discourse should not mention *Red* leaving the *Manor* or similarly *Blue* arriving and being angry at *White*. Conveying this information too soon would cause an undesired inference that would violate the causal link corresponding to the `(not (blv (has-param sm murderer B)))` invariant of the decomposition frame corresponding to the `pretend (occurs sm (murder R W M))` abstract action. This mechanism ensures that causal links are preserved even when the threat does not come directly from the effects of a step.

This prototype uses a C# implementation of the DPOCL algorithm extended to support the generation of inferences. The planner takes as input a story plan in a PDDL-like format and creates a discourse problem composed of the reified story and a set of discourse goals also in PDDL. The planning domain is composed from two sets of operator libraries written in PDDL: a set of discourse operators and a set of narration operators and idioms that, in the case of our first prototype, are specific to text. The planner uses the domain, problem, and a set of inference rules to compute a plan that is then transformed to text via a natural language realizer.

Conclusion and Future Work

We have presented in this paper a plan-based tripartite representation of narrative. This representation enables reasoning about audience inferences, whether to achieve discourse goals or to avoid thwarting narrative effects that require the audience to not hold specific beliefs. This model separates discourse actions, which have intended effects on the mental model of the audience, from narration actions, which correspond to the surface realization of these actions in specific media. This separation allows for the reusability of medium-independent discourse knowledge while maintaining reasoning capabilities on the expressive power of different media.

The knowledge representation and narrative generation process presented in this paper provide a framework for medium-independent and domain-independent discourse reasoning. However, the reasoning itself still relies strongly on *ad hoc* expert knowledge encoded as handcrafted planning operators. Similarly, the inferences rules used for the generation of the example presented are domain-dependent hand-written rules, but the representation could account for more complex inference mechanisms. Future work will use this framework as a base to operationalize cognitive models of narrative understanding (e.g., (Cardona-Rivera et al. 2012), (Niehaus and Young 2014)), and integrate them into generative processes. Moreover, while it takes into account the inter-dependency between discourse and narration choices, this representation still enforces a pipeline model between story and discourse generation, and between the choice of narration actions and their realization, as pointed out by Ronfard and Szilas (2014). Future work will move toward a more unified model of narrative generation.

Acknowledgements

This material is based upon work supported in whole or in part with funding from the Laboratory for Analytic Sciences (LAS). Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the LAS and/or any agency or entity of the United States Government.

References

Appelt, D. E. 1985. *Planning English Sentences*. Cambridge: Cambridge University Press.

Bae, B.-C., and Young, R. 2014. A Computational Model of Narrative Generation for Surprise Arousal. *IEEE Transactions on Computational Intelligence and AI in Games* 6(2):131–143.

Callaway, C. B., and Lester, J. C. 2002. Narrative prose generation. *Artificial Intelligence* 139(2):213–252.

Cardona-Rivera, R. E.; Cassell, B. A.; Ware, S. G.; and Young, R. M. 2012. Indexter: A computational model of the event-indexing situation model for characterizing narratives. In *The Workshop on Computational Models of Narrative at the Language Resources and Evaluation Conference*, 32–41.

Chatman, S. B. 1980. *Story and discourse: narrative structure in fiction and film*. Ithaca, N.Y.: Cornell University Press.

Cheong, Y., and Young, R. 2014. Suspenser: A Story Generation System for Suspense. *IEEE Transactions on Computational Intelligence and AI in Games*.

Cohen, P. R., and Perrault, C. R. 1979. Elements of a plan-based theory of speech acts. *Cognitive Science* 3(3):177–212.

Elson, D. K., and Mckeown, K. R. 2007. A platform for symbolically encoding human narratives. In *In Proceedings of the AAAI Fall Symposium on Intelligent Narrative Technologies*.

Ferguson, G. M. 1992. Explicit Representation of Events, Actions and Plans for Assumption-Based Plan Reasoning. Technical report, University of Rochester, Rochester, NY, USA.

Genette, G. 1980. *Narrative Discourse: An Essay in Method*. Cornell University Press.

Grosz, B. J., and Sidner, C. L. 1986. Attention, intentions, and the structure of discourse. *Computational linguistics* 12(3):175–204.

Hoek, M.; Theune, M.; and Linssen, J. 2014. Generating Game Narratives with Focalization and Flashbacks. In *Proceedings of 3rd Workshop on Games and NLP (GAMNLP-14) held at AIIDE-14*.

Hovy, E. H. 1990. Pragmatics and Natural Language Generation. *Artificial Intelligence* 43(2):153–198.

Jhala, A., and Young, R. M. 2010. Cinematic Visual Discourse: Representation, Generation, and Evaluation. *IEEE Transactions on Computational Intelligence and AI in Games* 2(2):69–81.

Lambert, L., and Carberry, S. 1991. A tripartite plan-based model of dialogue. In *Proceedings of the 29th annual meeting on Association for Computational Linguistics*, 47–54. Association for Computational Linguistics.

Mann, W. C., and Thompson, S. A. 1988. Rhetorical structure theory: Toward a functional theory of text organization. *Text* 8(3):243–281.

Maybury, M. T. 1991. Planning Multimedia Explanations Using Communicative Acts. In *Proceedings of the Ninth National Conference on Artificial Intelligence - Volume 1*, AAAI'91, 61–66. Anaheim, California: AAAI Press.

Maybury, M. T. 1992. Communicative acts for explanation generation. *International Journal of Man-Machine Studies* 37(2):135–172.

Montfort, N. 2011. Curveship's Automatic Narrative Style. In *Proceedings of the 6th International Conference on Foundations of Digital Games*, FDG '11, 211–218. New York, NY, USA: ACM.

Moore, J. D., and Paris, C. L. 1993. Planning text for advisory dialogues: Capturing intentional and rhetorical information. *Computational linguistics* 19(4):651–694.

Niehaus, J., and Young, R. M. 2014. Cognitive models of discourse comprehension for narrative generation. *Literary and Linguistic Computing* 29(4):561–582.

Pratt, M. L. 1977. *Toward a Speech Act Theory of Literary Discourse*. Bloomington: Indiana Univ Pr.

Riedl, M., and Young, R. 2010. Narrative Planning: Balancing Plot and Character. *Journal of Artificial Intelligence Research*.

Riedl, M. O.; Stern, A.; Dini, D.; and Alderman, J. 2008. Dynamic experience management in virtual worlds for entertainment, education, and training. *International Transactions on Systems Science and Applications, Special Issue on Agent Based Systems for Human Learning* 4(2):23–42.

Rishes, E.; Lukin, S. M.; Elson, D. K.; and Walker, M. A. 2013. Generating Different Story Tellings from Semantic Representations of Narrative. In Koenitz, H.; Sezen, T. I.; Ferri, G.; Haahr, M.; Sezen, D.; and Catak, G., eds., *Interactive Storytelling*, number 8230 in Lecture Notes in Computer Science. Springer International Publishing. 192–204.

Ronfard, R., and Szilas, N. 2014. Where story and media meet: computer generation of narrative discourse. In *Computational Models of Narrative*.

Swartjes, I., and Theune, M. 2006. A Fabula Model for Emergent Narrative. In Gbel, S.; Malkewitz, R.; and Iurgel, I., eds., *Technologies for Interactive Digital Storytelling and Entertainment*, volume 4326 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg. 49–60.

Theune, M.; Slabbers, N.; and Hielkema, F. 2007. The Narrator: NLG for digital storytelling. In *Proceedings of the Eleventh European Workshop on Natural Language Generation*, 109–112. Association for Computational Linguistics.

Young, R. M., and Moore, J. D. 1994. DPOCL: A principled approach to discourse planning. In *Proceedings of the Seventh International Workshop on Natural Language Generation*, 13–20.

Young, R. M.; Pollack, M. E.; and Moore, J. D. 1994. Decomposition and Causality in Partial-order Planning. In *AIPS*, 188–194.

Young, R. M. 2007. Story and discourse: A bipartite model of narrative generation in virtual worlds. *Interaction Studies* 8(2).