

Generating Plotlines about Attempting to Avert Disasters

Paper type: Short Paper

Pablo Gervás

Facultad de Informática
Universidad Complutense de Madrid
Madrid, 28040, Spain
pgervas@ucm.es

Abstract

Narratives about disasters triggered by a chain of connected incidents need to combine very diverse elements, such as the set of possible problematic events, the causal relations between these events and their potential consequences, the set of solutions that might be applied, the relations of the solutions to the problems and the expected effects of the solutions. Successful modeling of the search space of such elements would provide means for generating plot lines based on attempts to avert disasters. A prototype has been constructed that combines features to model causality, set sequences of actions to mirror emergency protocols, probabilistic information on event sequencing and timed injection of events. The relative merits of the prototype are discussed and conclusions concerning the general task of modeling this type of narrative are drawn.

Introduction

Plot lines where a small incident sets off a chain of events that can potentially lead to catastrophic consequences are a staple ingredient of entertainment fiction. They are usually combined with additional plot lines outlining attempts to break the chain of causality leading to disaster. Of these attempts, all but the very last one will usually fail, taking the story to the brink of disaster before success is achieved.

In order to obtain a computational model that captures the characteristics of the search space with sufficient detail, the following features need to be modeled: 1) interest (events corresponding to normal/acceptable operation are not of interest to the narrative), 2) causality (of the kind that captures the relationship between certain events and their consequences), 3) projection of current events forward into time (in order to foresee unwanted developments and plan against them) 4) potential solutions (that might be applied to stop the problems).

The present paper purposefully sets out to develop a model designed to capture these aspects of the problem. Although we assume that accurate modeling of all the aspects is beyond the scope of a simple paper, we will consider an initial approximation capable of creating narratives that are both realistic and interesting.

Previous Work

A number of aspects are relevant to the modeling of these types of narrative: their relation to simulations of the physical systems involved, their adversarial nature – disaster waiting to happen vs. characters attempting to avert it – and the role of causality in them.

The construction of computer simulations of nuclear power plants as physical models (Lu 1999) has been exploited as a platform for inexpensive cognitive psychology research (Ulrich et al. 2017). However, most of the events in such simulations are likely to correspond to regular performance of the plant, and unlikely to merit inclusion in a narrative that aims to be interesting. A model aimed at generating interesting narratives needs to focus on events that fall outside the regular performance of the plant.

The task of responding to emergencies is reactive: initiating incidents trigger chains of events that put the plant and emergency responders set in motion plans to avert disaster. This opposition has been modeled in computational narrative before (Dehn 1989) in terms of authors aiming to thwart some of the goals of their characters as a means to create interest in the plot.

Models of causality based on defining preconditions for actions to represent causal relations have been used successfully to model narrative in terms of planning problems (Riedl 2004). An initial situation and a desired goal are provided and the planner finds a sequence of events to connect them. Adversarial conditions are considered in interactive narrative (Porteous, Cavazza, and Charles 2010), where an initial plan by the system may be altered by the actions of the user and the system replans periodically to ensure its final goal is always achieved.

Generating Plotlines based on Averting an Imminent Disaster

The present paper aims to develop a procedure for constructing plot lines built around the idea of an upcoming disaster set in motion by a small incident, and a series of attempts to prevent it. To develop a system capable of modelling this type of behaviour, the following assumptions are made: 1) uneventful operation of the system warrants no narrative, 2) some kind of triggering incident (that constitutes a problem or has the potential to lead to a problem) sets a narrative in

motion 3) a sequence of events linked by causality is set in motion, such that 4) (at least) the final event in that sequence constitutes a problem; 5) the problem triggers a (number of) reactions, and 6) the reactions tend to be plans (possibly more than one) all designed to solve the problem (or some part of it) in some way, 7) reactions may themselves trigger additional causal chains of events and 8) further incidents (also problematic) may take place that break the causal chains arising from reactions.

Disaster movies usually combine a large number of types of events (meteorology, natural disasters, human error, terrorism, mismanagement...). Exhaustive modelling of all these types of events is beyond the scope of this paper. In order to provide a relatively simple domain with the desired complexity where causality relations between events can be defined in objective terms an existing knowledge-base for accidents in nuclear power stations¹ is chosen as underlying world model.

Traditional planning systems are ill suited for modelling this type of problem because: 1) the model cannot be deterministic (problems may or may not happen in different runs of the system, and these changes in behaviour need to be achieved without having to change the rules that govern the causality of the system), 2) emergency protocols take the form of timed sequences of events (responders are not necessarily aware of the causal chains that they are intended to trigger), and 3) potential chains of causality initially identified by the system may need to be broken by further problematic events (this allows the capture of mismatches between expectations and real system reactions).

Customised Representation of Narrative

Events are represented as atomic propositions. The causal chaining of events is modeled by the possibility of assigning to each event a set of *preconditions* that may trigger it (the activation of the preconditions is said to cause the following events). A description of the lexicon entries for events concerning how heat affects people are listed in Table 1, together with an example of how they would be chained by the system into a causal sequences of events (in the absence of interventions in the form of emergency responses).

Time is modelled in a step fashion, with a succession of turns taking place, and a set of events taking place simultaneously at each turn. The set of events that take place in a given turn are said to be *activated* on that turn.

An initial requirement is that the system allow for both user-driven *configurable* mode – to allow narratives to be generated on demand for specific situations – and an *autonomous* generation mode – where the system generates narratives by exploring the search space of possible narratives with no guidance from the user. To allow for this possibility, the system includes functionality for injecting events at particular moments in time. An *injection schedule* indicates the relative timing for a given set of events, starting from the point at which the injection is triggered. Table 2 shows an example of an injection schedule.

¹Currently under development for the ADARVE project, ref. SUBV-20/2021, funded by the Spanish Nuclear Security Council.

The second basic requirement is that the system be capable of representing the occurrence of problems that may lead to consequences, most of them undesirable. The occurrence of problems may be represented using the functionality for injecting events.

Whenever an event is activated, the system uses causal chaining based on preconditions to identify any events that may be triggered as consequences. The full chain of consequences of an event is computed during the turn when the event is activated, but they do not start to be activated until the following turn. At each turn, the next level of consequences from preceding events is activated unless some solution has blocked them.

Table 3 presents a set of causal chains produced by the system that capture various possible problems modelled for nuclear power plants and the outcomes predicted if no action is taken.

Modelling Emergency Response to Problems

To model responses to problems, the system contemplates a set of elaborate responses (emergency plans) to problematic events (these elaborate responses are represented as patterns of actions to be undertaken in particular order and following an established relative timing between them).

Emergency responses are encoded in the system at three different levels. At the first level, a solution for a given problem associates the problem itself with the name of a particular plan of actions. At the second level, the actual plans to be carried out are represented as a timed sequence of events to be injected into the system. These timed sequences are also represented as injection schedules. At the third level, the system models consequences of actions taken as part of plans in the same way that it modelled consequences of problems. This allows these emergency plans to be expanded into causal chains triggered by actions in plans without having to explicitly list the consequences in the description of the plan. Causal links between plan actions and further events are modelled in the lexicon. Table 4 shows an example of a causal chain arising from a plan. This example shows how the problem chain shown in Table 1 may be interrupted by the application of an emergency plan.

Probability Driven Construction of Event Consequences

Causal relations between events are not always completely deterministic. Often consequences follow the events that trigger them only in some cases. This peculiarity should also be included in the model.

The system considers that, for each event that is activated, the set of its possible consequences needs to be compiled. For the events in this set, a conditional probability of occurrence given the triggering event must be considered. To inform this process, a *probability of occurrence* is associated with each event. For a more realistic model, the conditional probabilities of each consequence given its trigger should be considered, but a single probability is considered an acceptable approximation for an easier initial model.

At each turn, the consequences of events activated in the preceding turn are considered for activation. Based on

Trigger	Event
(injected event)	HeatReachesPeople
HeatStartAffectingPeople CAUSE [HeatReachesPeople]	HeatStartAffectingPeople
PeopleSufferFromHeat CAUSE [HeatStartAffectingPeople]	PeopleSufferFromHeat
PassOutFromHeat CAUSE [PeopleSufferFromHeat]	PassOutFromHeat
PeopleDie CAUSE [PassOutFromHeat,SufferTerminalRadiationPoisoning]	PeopleDie

Table 1: Extract of lexicon entries for events concerned with the effect of heat on people (first column), together with an example causal sequence produce by the system.

Time offset	Event
0	Tornado
1	DamageToGenerator
2	NoFuelForDieselPoweredElements

Table 2: Example of injection schedule showing a tornado, damage to a generator, and lack of fuel.

DamageToDieselPumps
DieselPumpsNotWorking
NuclearReactorStartsOverheating
OverheatsNuclearReactor
NuclearReactorReachesDangerousTemperature
CoolantEvaporates
HeatReachesPeople
NuclearReactorStartsToMelt
RadioactiveMaterialExposed
HeatStartAffectingPeople
NuclearReactorMeltdown
PeopleSufferFromHeat
PassOutFromHeat
PeopleDie

Table 3: Examples of causal chains for problems in a nuclear power plants and the outcomes predicted if no action is taken. Horizontal lines represent turn transitions.

Problem/HeatReachesPeople
Problem/HeatStartAffectingPeople
StartingPlanToSolve/HeatReachesPeople
Solution/RemovePeopleToSafety
Problem/PeopleSufferFromHeat
Solution/PeopleSafe
PlanSucceeded@HeatReachesPeople

Table 4: Causal chain arising from the activation of an emergency plan.

a random factor, consequences in this set are activated if a randomly generated number falls under the probability threshold associate to the event.

There is another aspect of consequence that needs to be captured in the model. In certain cases, two contrasting potential consequences of the same event are registered, and only one of them should be activated in any particular situation (*disjunctive branching*). To capture this feature, situations of this type are also explicitly encoded in the system and a probability is assigned to them that is used by the system to select only one of them when expanding.

The introduction of probabilistic information becomes a useful tool for customising system performance to meet the different required modes. For the *autonomous* mode, weighted random choice informed by the recorded probabilities allows for consistent variation in the outcomes. Explicit customisation of the relative values of the probabilities allows the user to constrain the search for narratives to particular subsets of the search space, allowing operation in the *configurable* mode.

Disjunctive branching is used to capture the dichotomy between success and failure of actions undertaken in pursuit of particular goals in response to emergencies. In this way, when the system is being run in the *autonomous* mode, plans may succeed or fail regardless of diligent application of established protocols, making the narratives more interesting. For the *configurable* mode, the user may tailor the probabilities used for disjunctive branches to drive system outcomes to particular alternatives.

Compilation of a Full Plot

The system operates on an injection schedule taken as input, that determines a certain sequence of events that happen at particular times. These are the incidents that create the cascade of problems. The system progressively compiles chains of consequences for events at a given turn and determines which events in those chains will happen at the next turn. It also compiles which responses may be undertaken to break the chains of undesirable consequences. Table 5 shows an example of a plot line generated for a combination of a damaged transformer and lack of fuel for diesel generators.

Due to its reliance of probabilities to drive the final outcome, subsequent runs of the system will produce plots outlines for different narratives for a given input. Additionally, the probabilities may be tailored The system can be used as a co-creation assistant, allowing the user to compile a set of plot outlines for a given configuration (input + defined probabilities) or even for combinations of different inputs and values for the probabilities.

State	
0	Problem/DamageToPowerTransformer (injected)
1	Problem/LowPowerToOperatePlant (from <i>DamageToPowerTransformer</i>)
2	StartingPlanToSolve/LowPowerToOperatePlant Solution/ShutDownNon-vitalSystems ToSolve/@LowPowerToOperatePlant PlanEnded ToSolve/@LowPowerToOperatePlant
3	Solution/MorePowerAvailableForVitalSystems (from <i>ShutDownNon-vitalSystems</i>) Problem/ElectricPumpsShutDown (from <i>ShutDownNon-vitalSystems</i>)
4	PlanSucceeded ToSolve/LowPowerToOperatePlant StartingPlanToSolve/ElectricPumpsShutDown Solution/HookUpDieselGeneratorToElectricPumps ToSolve/@ElectricPumpsShutDown PlanEnded ToSolve/@ElectricPumpsShutDown
5	Problem/OverheatsNuclearReactor (from <i>ElectricPumpsShutDown</i>) Problem/NoFuelForDieselPoweredElements (injected)
6	Problem/NuclearReactorReachesDangerousTemperature (from <i>OverheatsNuclearReactor</i>) Problem/CoolantEvaporates (from <i>OverheatsNuclearReactor</i>) Problem/DieselPumpsNotWorking (from <i>NoFuelForDieselPoweredElements</i>) StartingPlanToSolve/NoFuelForDieselPoweredElements Solution/FuelDelivery ToSolve/@NoFuelForDieselPoweredElements PlanEnded ToSolve/@NoFuelForDieselPoweredElements
7	Problem/HeatReachesPeople (from <i>NuclearReactorReachesDangerousTemperature</i>) Problem/NuclearReactorStartsToMelt (from <i>NuclearReactorReachesDangerousTemperature</i>) Problem/RadioactiveMaterialExposed (from <i>CoolantEvaporates</i>) Solution/StartsDieselGenerator (from <i>Solution/FuelDelivery</i>)
8	PlanSucceeded ToSolve/NoFuelForDieselPoweredElements Problem/HeatStartAffectingPeople (from <i>HeatReachesPeople</i>) Problem/NuclearFuelMeltingPointReached Solution/StartsElectricPump (from <i>StartsDieselGenerator</i>) StartingPlanToSolve/HeatReachesPeople Solution/RemovePeopleToSafety ToSolve/@HeatReachesPeople PlanEnded ToSolve/@HeatReachesPeople StartingPlanToSolve/RadioactiveMaterialExposed Solution/DivertReactorCoolantToWastePool ToSolve/@RadioactiveMaterialExposed PlanEnded ToSolve/@RadioactiveMaterialExposed
9	PlanSucceeded ToSolve/@ElectricPumpsShutDown Solution/PeopleSafe (from <i>RemovePeopleToSafety</i>) PlanSucceeded ToSolve/@HeatReachesPeople Solution/FillsUpWastePool (from <i>DivertReactorCoolantToWastePool</i>) PlanSucceeded ToSolve/@RadioactiveMaterialExposed

Table 5: Example of Generated Plot. Injected events are shown in **bold**, the relations of causality between the various elements in *italic*, the responses in Typewriter font

Discussion

The prototype includes various techniques to allow it to capture the full set of requirements identified from the formative analysis of the desired type of narratives.

The procedure for construction of consequence trees from an initiating event (problem) mirrors the representation of causal relations as used in planning-based representation of narrative, but it differs from them in that the chaining applied is not goal-driven. This is because the behaviour being modelled does not correspond to intentional actions, but rather to expanding consequences of given events. As in planning solutions applied to interactive narrative, only part of each plan built is actually used as a contribution of the narrative, because events from attempted solutions may cut off consequences of problems and further incidents may block some of the solutions.

The introduction of probabilistic information to drive the process of expanding events into consequence trees allows selective exploration of the search space. Customisation of the set of probabilities for specific runs allows relative steering of the outcome narratives.

The mechanism for injecting particular events according to a fixed schedule allows configuration of the system to produce narratives triggered by specific events, or modified at particular points by specific events.

The mechanism for modelling emergency response to problems in terms of set sequences of relatively scheduled events allows the representation of known emergency proto-

cols or accepted plans of action in each particular domain.

All these mechanisms may be useful in the context of larger storytelling systems, where a disaster-averting plot line may need to be tailored to combine well with other plot lines (romantic interest, social problem, rivalry between characters...). Plot lines of other types may be generated by completely different mechanisms (Gervás 2021), created by hand or reused from known plot schemas (Concepción, Gervás, and Méndez 2016). Automated combination of more than one plot line may be considered (Gervás, Concepción, and Méndez 2022).

The solution employed in this paper for knowledge representation is based on simple propositional atoms, each one representing an event. This solution is considerably clumsier than those used in existing attempts at modelling narrative, such as those based on planning, grammar, or case based reasoning. However, it presents the important advantage of being an acceptable approximation to all the formalisms employed in each of those paradigms. Since the problem being considered in this paper appears to require a combination of several of these techniques, it was important to allow a solution compatible with all. Refinement of the knowledge representation can be considered as further work. In that process of refinement, the specific features that characterise each of the paradigms may be considered as possible extensions for the model, but in all cases they must be evaluated for compatibility with the elements from other paradigms that have been identified as necessary.

Conclusions

The model presented in this paper captures adequately the requirements identified for narratives about averting disasters triggered as final consequences of some initial incident.

The methodology followed in the paper, basing the construction of the model on combining a set of technologies to ensure that it contemplates the main features of the problem domain, places the priority in finding an adequate solution to a real problem. This is compatible with existing theoretical analyses of the complexity of narrative that suggest that successful modeling of narrative requires inclusion of specific representational solution for the many different aspects (Gervás and León 2014) that contribute to the richness of the medium.

With respect to future work, a number of potential lines are considered. First, the refinement of the knowledge representation schema to include information on the agents participating in the events would lead to a finer grained modeling of both the problems and the solutions being considered in the narratives. Second, since the successful modeling of the narratives in the inspiring set has been shown to require features of causality, of set sequences of events, and of probabilistic information, refinements to the model may be considered based on planning, case-based reasoning, and Bayesian inference networks.

Overall, the proposed solution remains a simple initial approximation to the problem, but it serves to highlight the need to contemplate a number of techniques to capture the full spectrum of features that are found in the narratives in the inspiring set.

Author Contributions

Pablo Gervás ideated and wrote the paper alone.

Acknowledgments

This paper has been partially supported by the CANTOR project (PID2019-108927RB-I00) funded by the Spanish Ministry of Science and Innovation and the the ADARVE (Análisis de Datos de Realidad Virtual para Emergencias Radiológicas) project (SUBV-20/2021) funded by the Spanish Nuclear Security Council (Consejo de Seguridad Nuclear, CSN). Some of the ideas in the paper benefited from discussions with Carlos León, Gonzalo Méndez and Eugenio Concepción.

References

- Concepción, E.; Gervás, P.; and Méndez, G. 2016. Mining knowledge in storytelling systems for narrative generation. In *CC-NLG: Computational Creativity in Natural Language Generation (@INLG2016)*. Edimburgh, UK: ACL Anthology.
- Dehn, N. 1989. *Computer story-writing: the role of reconstructive and dynamic memory*. Ph.D. Dissertation, Yale University.
- Gervás, P., and León, C. 2014. The need for multi-aspectual representation of narratives in modelling their creative process. In *2014 Workshop on Computational Models of Narrative*. Quebec City, Canada: Scholoss Dagstuhl OpenAccess Series in Informatics (OASIS).
- Gervás, P.; Concepción, E.; and Méndez, G. 2022. Evolutionary construction of stories that combine several plot lines. In *Computational Intelligence in Music, Sound, Art and Design – 11th International Conference, EvoMUSART 2022*. Madrid, Spain: Springer.
- Gervás, P. 2021. Computational models of narrative creativity. In Machado, P.; Romero, J.; and Greenfield, G., eds., *Artificial Intelligence and the Arts: Computational Creativity, Artistic Behavior, and Tools for Creatives*. Cham: Springer International Publishing. chapter 9, 209–255.
- Lu, S. 1999. Dynamic modelling and simulation of power plant systems. *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy* 213(1):7–22.
- Porteous, J.; Cavazza, M.; and Charles, F. 2010. Applying planning to interactive storytelling: Narrative control using state constraints. *ACM Trans. Intell. Syst. Technol.* 1(2).
- Riedl, M. 2004. *Narrative Planning: Balancing Plot and Character*. Ph.D. Dissertation, Department of Computer Science, North Carolina State University.
- Ulrich, T. A.; Lew, R.; Werner, S.; and Boring, R. L. 2017. Rancor: A gamified microworld nuclear power plant simulation for engineering psychology research and process control applications. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 61(1):398–402.