



UvA-DARE (Digital Academic Repository)

Legal Knowledge Conveyed by Narratives: Towards a Representational Model

Sileno, G.; Boer, A.; van Engers, T.

DOI

[10.4230/OASlcs.CMN.2014.182](https://doi.org/10.4230/OASlcs.CMN.2014.182)

Publication date

2014

Document Version

Final published version

Published in

5th Workshop on Computational Models of Narrative

License

CC BY

[Link to publication](#)

Citation for published version (APA):

Sileno, G., Boer, A., & van Engers, T. (2014). Legal Knowledge Conveyed by Narratives: Towards a Representational Model. In M. A. Finlayson, J. C. Meister, & E. C. Bruneau (Eds.), *5th Workshop on Computational Models of Narrative: CMN'14, July 31–August 2, 2014, Quebec City, Canada* (pp. 182-191). (Open Access Series in Informatics; Vol. 41). Schloss Dagstuhl - Leibniz-Zentrum für Informatik. <https://doi.org/10.4230/OASlcs.CMN.2014.182>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Legal Knowledge Conveyed by Narratives: Towards a Representational Model

Giovanni Sileno, Alexander Boer, and Tom van Engers

Leibniz Center for Law, University of Amsterdam
The Netherlands
g.sileno@uva.nl, aboer@uva.nl, vanengers@uva.nl

Abstract

The paper investigates a representational model for narratives, aiming to facilitate the acquisition of the systematic core of stories concerning legal cases, i.e. the set of causal and temporal relationships that govern the world in which the narrated scenario takes place. At the discourse level, we consider narratives as sequences of *messages* collected in an *observation*, including descriptions of agents, of agents' behaviour and of *mechanisms* relative to physical, mental and institutional domains. At the content level, stories correspond to synchronizations of embodied *agent-roles* scripts. Following this approach, the *Pierson v Post* case is analyzed in detail and represented as a Petri net.

1998 ACM Subject Classification H.1.2 Human information processing, I.2 Artificial Intelligence

Keywords and phrases story representation, story acquisition, legal narratives, knowledge representation, agent-roles, causation, expectations, agent-based modeling, petri nets

Digital Object Identifier 10.4230/OASICS.CMN.2014.182

1 Introduction

Legal activity provides excellent examples of the operational use of narratives, for instance in *adjudication*. The interaction between parties, witnesses, experts and judges includes narrative acts. Moreover, if the case introduces a precedent, the publication of its proceedings may be seen in itself as a narrative act, meant to inform the legal system of novelties concerning social interactions and their legal interpretation. This phenomenon reproduces at a systemic scale what happens in the daily legal practice, as legal experts usually rely on prototypical or hypothetical cases when they explain or unravel a certain legal problem.

Our objective is therefore to investigate an adequate representational model for cases (historical, hypothetical, etc.). At first sight, the domain of application is a specific class of narratives, but, in reality, it is a structural component of all narratives, related to the socio-institutional interpretation of behaviour; the approach we propose may be plausibly used as well on folk tales, mythology, etc. if the motive is to investigate possible underlying normative indications.

Background. Several models have been proposed in the literature in order to represent *stories*, introducing concepts like *story grammars* [21], *scripts* [1], *plot units* [13], *multi-level representations* [17], *doxastic preferences* [15], *story intention graphs* [9] etc. All these contributions target structural components of narratives, primarily with the intent of reproducing the in-depth knowledge mechanisms specifically behind story understanding, story generation and summarization. Although the present work can be seen as a follow-up of these contributions, our focus is slightly different.



© Giovanni Sileno, Alexander Boer, and Tom van Engers;
licensed under Creative Commons License CC-BY

5th Workshop on Computational Models of Narrative (CMN'14).

Editors: Mark A. Finlayson, Jan Christoph Meister, and Emile G. Bruneau; pp. 182–191

OpenAccess Series in Informatics



OASICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

We are primarily targeting the knowledge involved in a certain (legal) case, taking into account the possibility of receiving different explanations of a given sequence of events, usually from different point of views, allowing to consider alternative interpretations of agents' (legal) positions and intentions. Assuming that a *story* exists in the mind of the narrator/listener, we are mostly concerned by the problem of *story acquisition*, in the sense of acquiring the interpreted content in a formal representation, eventually supported by specific diagramming tools. Rather than text annotation practices, our elicitation model has more similarity with *scenario-based modeling*, used in software engineering, e.g. [11]. This paper in particular focuses on Petri nets.¹

Knowledge acquisition is inherently coupled with finding the right representational model for the target domain. As Lwe observes in [16], for instance, *plot units* can represent causality, but not expectation; the *doxastic preference framework*, conversely, models the second, but not the first. As we need both these features to adequately characterize social interactions, this work proposes a possible solution.

The paper proceeds as follows. The theoretical framework is introduced incrementally, starting with the narrative model (2) and with a case example of narrative. We define then the story model (3), further refined with agent-role models (4). To conclude, we briefly present the representation of the case as a Petri net (5). Discussion ends the paper.

2 The narrative model

Any *narrative* always manifests three ontological layers, present at the same time: the **discourse**, the **story** and the **conversation**.² They are respectively signal, meaning, and relevant components of the social context—like knowledge/intents of narrator and listener—that concur to the generation/interpretation of such act of communication. The “same” story can be reported with different discourses, because the discourse defines the order and the form (verbal and non-verbal) in which the content is provided. The specific choices of transformation from story to discourse (in generation) and from discourse to story (in interpretation) are influenced by the conversation layer.

The narrative content received by an interpreter is in the form of a sequence of *messages*, which essentially correspond to *speech acts* [22]. The complete set of messages is called *observation* and constitutes a superficial layer of meaning, reporting events, states, and possibly explicit dependencies between them, i.e. the *foreground mechanisms*. In general, the interpreter has to integrate this content with other *background* components in order to complete the model with dependencies and facts that are missing in the narration, but he recognizes as relevant to explain its occurrence.³

Narratives often contain characters that *tell* something. In respect to these quoted or indirect speeches, which possibly constitute *nested narratives*, the higher-level story (brought by the narrator) may include parts of their discourse and clues about their conversation layers (e.g. position, knowledge, intentions of the participants).

¹ From a wider perspective, however, we target an *integrated development environment* (IDE) accounting multiple views, and allowing an incremental refinement of the elicited content. We give a preliminary example in [25], referring to UML diagrams (Message Sequence Charts), topology diagrams, Petri nets and scripts in AgentSpeak(L), a logic programming language for cognitive agents.

² The introduction of *story/discourse* distinction is usually associated to the Structuralists (Barthes, Todorov, Genette, etc.). Before them, the Russian formalists (Propp, Shklovsky) used the terms *fabula/syuzhet*. For the use of the term *conversation*, see for instance Young in [27]; other authors prefer *narration*.

³ Based on these concepts, we presented in [26] a preliminary implementation of the process of interpretation, constructed on *explanation-based argumentation*.

► **Example 1.** As example of narrative, we consider a well known case in Property Law and in AI & Law [6]: *Pierson v Post* (1805)⁴. The story is basically the following:

Post was hunting a fox with a horse and hounds in a wild and uninhabited land, and was about to catch it, but Pierson, although conscious of Post's pursuit, intercepted, killed and took the animal.

Both claimed the fox, the first appeal had found for Post, but this court reverted the previous result. The different positions are expressed by two judges: Tompkins (*majority*) and Livingston (*dissent*). The first, supported by classical jurisprudence, claims that:

Tompkins: Possession of a *fera naturae* occurs only if there is occupancy, i.e. taking physical possession. Pierson took the animal, so he owns it.

where *fera naturae* is an animal wild by nature. The second argues that:

Livingston: If someone starts and hunts a fox with hounds in a vast and uninhabited land has a right of taking the fox on any other person who saw he was pursuing it.

Both interpretations are relevant for our purposes, and create two different stories (also in a practical sense, as they would bring about different consequences). The two judges can be seen as playing the role of two different modelers, providing different mechanisms.⁵

3 The story model

In narratology the story layer is usually called the **fabula**: “a series of logically and chronologically related events [...]” [4]. This name dates back to Propp, which, altogether with the Russian formalists, started considering each event in the story as *functional*, i.e. a part of a whole sequence, necessary to bring the narrated world from initial conditions to a certain conclusion.

Following this perspective, as first definition, we may consider the story as a chain of events (a strictly ordered set):

$$\mathcal{E} = \{e_1, e_2, \dots, e_n\} \quad (1)$$

In addition, specific circumstances may be described in correspondence to the occurrence of an event. A more complete definition of story should consider the following chain:

$$C_0 \xrightarrow{e_1} C_1 \xrightarrow{e_2} \dots \xrightarrow{e_n} C_n \quad (2)$$

where e_i are associated to *transitions* and C_i is a set of conditions assumed to *continue* at least until the occurrence of e_i . Amongst those continuants we find also *existents*, as characters, objects, etc.

⁴ Source text: <http://www.facstaff.bucknell.edu/kinnaman/Piersonv.htm>.

⁵ In these terms, court proceedings can be seen as pushing to the public foreground institutional mechanisms not adequately defined in certain contexts. See [14] about the role of narratives in respect to *tacit* knowledge.

Plot. The definition of story given above seems quite simple, but the manifold relations between *consequence* (logical, causal) and *consecutiveness* (informed by time, ordering, verbal tense) are actually very delicate to assess. Furthermore, two different chronological coordinates coexist in a narrative: a *story-relative* time, i.e. when the event has occurred in the story, and a *discourse-relative* time, i.e. when that event has been reported/observed.

The concept of *plot* is relevant in this issue.⁶ Trying to better scope the problem, we identify three *levels of constraints* on the ordering of events via the following method, first presented in [25] but hereby slightly refined:

1. We identify the events and conditions explicitly expressed in the story.
2. For each event, we identify which conditions/events in the story are direct requirement and consequence of its occurrence. If necessary, we integrate them with external knowledge. The relations elicited in this way constitute the *dependencies* (causal, logical) and place some **strong constraints** on the ordering of events.
3. Time positions, durations and use of verbal tenses in the narration are usually meant to give some landmark to the reader/listener. They are described in absolute or relative terms. Once interpreted, they create a relation between events/conditions *contingent to the story*.⁷ The correspondent positioning constitutes the **medium constraints**.
4. When we do not have any other information, a possible sequence is suggested by the discourse-relative time, i.e. relative position in which the events are given. The consequent representational outcome is *contingent to the discourse* and provides the **weak constraints** on the ordering of events. They complement the ordering resulting from the medium constraints.

In the previous section, we presented the story as a strict ordered set of events. However, it is easy to object to such a strict determination: (a) dependencies can be associated to no-time-consuming processes (e.g. logical dependencies); (b) events may occur simultaneously, when triggered by parallel sub-systems (e.g. two agents acting independently), unless there is an explicit temporal determination. Consequently, we weaken the previous strict temporal constraints (from $e_{i+1} > e_i$ to $e_{i+1} \geq e_i$) at least in these two cases. With these modifications, the set of events \mathcal{E} defined in (1) becomes a *partially ordered set*.

Evidently, most of the strong constraints proposed in step 2 are relative to the reader, as only some of these dependencies are explicitly presented by the narrator. They respectively define the background and foreground **mechanisms** of the story-system. Conversely, the medium and weak constraints are always explicitly addressed in the narrative and can be objectively extracted as intrinsic component of the **observation**. They describe *contingent* relations, which are also *contextual* when they are part of the **synchronization**: i.e. the necessary alignment with the mechanisms for the story to occur.⁸ Suggested by the term control-flow in programming, we call the whole composition of constraints *story-flow*.

⁶ Unfortunately, there are conflicting accounts about its definition in the literature. For some authors the plot coincide with how things are presented, and therefore it practically equals the discourse. According to an older tradition, usually associated to Forster [10], the *story* properly said is only the chronological sequence of events, while the plot is the causal and logical structure connecting them.

⁷ Contingent to the story/discourse means that even if this specific story/discourse reports this ordering, *a priori*, there may be as well stories in which an alternative ordering holds.

⁸ Considered as a whole, however, constraints may be conflicting. When this occurs, it is a symptom of a problem with the mechanisms and/or with the observation. There might be dependencies which are missing, not valid or not acceptable in that specific context, or there may be faults in the timing or the nature of the reported events. Not addressed in this paper, this problem is one of the objects of our current research on *model-based diagnosis* and *justification* [25].

- **Example 2.** Applying the method on the “brute” story⁹ of *Pierson v Post*, we have:
- conditions and events: Post hunting (c_1), Post being in an uninhabited land (c_2), Post catching the fox (c_3), Pierson being conscious of Post’s pursuit (c_4), Pierson intercepting the fox (e_1), Pierson killing the fox (e_2), Pierson taking the fox (e_3)
 - strong constraints: $e_1 < e_2 < e_3$ (intercept, kill, take necessarily occur in this order)
 - medium constraints: $\{c_1, c_2, c_3\} \xrightarrow{e_1} \{\}$, $\{c_4\} \xrightarrow{e_1} \{\}$
 - weak constraints: $e_1 < e_2 < e_3$

Conditions like ‘Post hunting’ refer to ongoing actions which started in an undefined moment before the given story occurs. Note how the weak constraints reproduce the strong constraints: this is a natural tendency. In general, however, other events may occur in the story with no relevance for the causal mechanisms applied by the interpreter or with no specific ordering necessity.

Talking about mechanisms, the judges propose two institutional interpretations, which, added to the previous strong constraints, produce two alternative stories:

- $\{c_6\} \xrightarrow{e_3} \{c_7\}$: the fox being a wild animal (c_6), the event of Pierson taking the fox (e_3) makes Pierson owner of the animal (c_7);
- $\{c_1, c_2, c_4\} \rightarrow \{c_8\}$: Post hunting in an uninhabited land (c_1 and c_2), and Pierson conscious of his pursuit (c_4) gives Post an exclusive right to catch the prey (c_8), and, consequently, Pierson is not permitted to.

The last constraint involves only conditions, it is a logical dependency; it can be translated similarly to the others if we take into account some condition-activating event, for instance the last one (Pierson becoming aware of Post’s hunting).

4 The agent-role model

So far, the foundational structure of the story model consists of events, generic continuant entities and relations between them. Amongst continuant entities, characters are particularly important in narratives, as they connect events all throughout the story with their direct or indirect participation in actions.

Similarly to Propp, who, investigating Russian folk-tales [20], abstracted characters to *roles* defined by recurrent patterns of actions (the Villain, the Hero, etc.), we abstract agents to **agent-roles** (first presented in [8]). Agent-roles are *roles*, as they refer to prototypical patterns of behaviour (correspondent or building on top of roles defined by institutions), and they are *agents*, because their behaviour is described via cognitive and motivational components. In previous works [25, 24], we presented a multi-layered framework to be used for the characterization of agent-roles, summarized in Table 1.

Integrating some fundamental mechanisms inspired by BDI architectures (*beliefs-desires-intentions*) we relate the primary elements of each layer with catalyzers and with components of other layers, in order to construct a more in-depth representation of the strong constraints. The layers essentially reproduce the general story scheme used by Bex and Verheij in [7], and synthetizing in turn what was proposed by Pennington and Hastie in [19]:

$$Motive \rightarrow Goal \rightarrow Action \rightarrow Consequences$$

⁹ We are referring to Searle’s theory on institutions [22], which distinguishes *brute*, raw facts—in the sense of belonging to the world of experience—from *institutional* facts, i.e. facts which are meaningful within an institutional framework. Only some of the brute facts *count as* institutional fact, depending on constitutive rules.

■ **Table 1** Multi-layered framework to be used for the characterization of agent-roles.

	Primary elements	Catalyzer elements
Motivational layer	Motive	Motivation
Intentional layer	Intention	Affordance (perceived power)
Action layer	Action, Attempt	Disposition (actual power)
Signal layer	Message, Outcome	–

This scheme serves as a template to model the behaviour of each agent. In general, however, the complete decomposition is not fully expressed in a narrative, but retroactively occurs in the mind of the interpreter, helped by narrative clues and following his own conceptualization of the world.

In order to trace the relevant dependencies, we follow a methodology consisting in the elicitation of (a) *motives*: we identify the events which are reasons for action; (b) contextual *affordances*: we write down which conditions, in this story-world, are sufficient for an agent to consider to be successful in starting a course of actions aiming to a certain intent; in doing this, we reconstruct as well the course of actions as a hierarchical plan; (c) *dispositional rules*, modeling the agent-independent dependencies holding in the story-world. Obviously different modelers may provide alternative solutions. However, for our acquisition purposes, we take a neutral perspective toward the problem of deciding which are the “right” mechanisms representing the world: if the story-model, when executed, reproduces the messages given with the observation, it is *valid*.

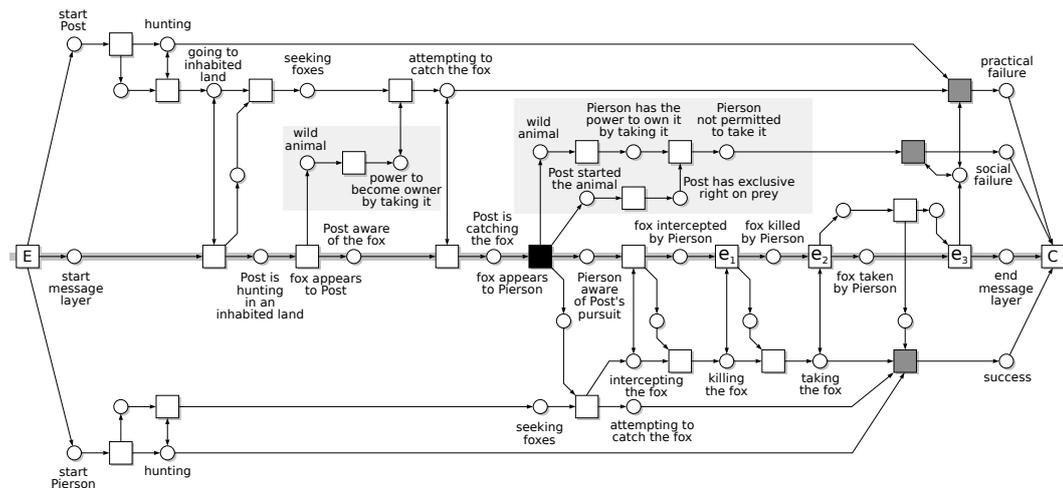
► **Example 3.** Executing these elicitation tasks on our story, a possible result is:

- for both Pierson and Post the core action is catching/taking the fox. A possible initiating motive is having seen the fox (not explicitly given), while the motivation would be associated to their involvement in a hunting activity or their proneness to hunt;
- the affordance of the core action, in respect to this story-world, is hidden to the reader. We assume that both of them thinks that they are able to take the fox in those conditions. The specific sequence of actions to be executed can be easily reconstructed via a chain of dependencies: e.g. you can take the fox, if you kill it; you can kill it, if you intercept it;
- if someone has already caught the fox, nobody can physically take it any more.

Institutional characterization. *Institutional power* (in legal terms also called *ability*) and *permission* can be described similarly to affordance, as they are assumed to be taken into account by the agent when he decides whether to proceed on a certain behaviour. The former is necessary for the action to be recognized by the social system; the latter identifies the liability of the agent to some enforcement action in case of violation.¹⁰ For the sake of brevity, we redirect the reader to [23] for further support on these claims.

Expectations, failures and successes. Relations as powers, obligations and permissions can be used to describe expectations about the social behaviour of the others in terms of normative possibility and necessity, just as affordances describe (expected) possibilities of action to agents in a certain environment, and dispositional rules describe the laws of a certain world. If the agent starts an action and is not successful in achieving his intent, then he acknowledges a *practical failure*. If the agent performs an act he considers institutionally

¹⁰ On the other side, *obligations* not yet fulfilled are prototypical motives for action for the addressee of the obligation. As permissions, they are generally associated to some form of enforcement.



■ **Figure 1** Simplified Petri Net representation of *Pierson v Post*.

meaningful, but this is eventually not recognized as such, he has still failed in respect to his intents. If the agent intentionally performs an act which is not permitted or does not satisfy a given duty, there is an *institutional failure*, but if the intended outcome is still successfully established, then it will still be a *practical success* for him, even if enforcement actions may reduce the general pay-off. The interactions between intents, practical/social expectations and actual outcome entail the failures or successes for the agents, also in social terms.

5 Translating the story-flow as a Petri Net

Within the formalisms used in computation, Petri nets¹¹ are one of the best established tools for *process modeling* and *analysis*. They mirror the definition of story model we constructed above, because they allow to explicitly divide conditions (represented with places) from events (with transitions), and reproduce changes of state via the movement of tokens, respecting the partial ordering property we associated to the story-model.

We have reported in Fig. 1 a simplified representation of our case story, obtained from the composition of its constraints, in order to informally explain the construction principles of our proposal. Three macro-areas can be recognized. The first, the *message layer*, corresponds to the central line, and contains the events/conditions provided by the observation. The arcs on this line chronologically reproduce the *synchronization* between external and internal events. Agents or other parallel subsystems (in our case, Pierson and Post) start together with the message layer and interact through it.¹² Their subnets contain elements belonging to all other three layers described in 4.

¹¹ A Petri net is a directed, bipartite graph with two types of nodes: *places* (visualized as circles) and *transitions* (bars). A place is connected only to transitions and vice-versa. One or more *tokens* (black dots) can reside in each place, while transitions can be *fired*, moving those tokens from their input places to their output places. Petri nets furnish a direct visual representation of the causal structure (via the network) and of the behaviour (via the movement of tokens) of the system.

¹² The picture shows some fundamental patterns. For example, a simple “writing” pattern is attached to “going to inhabited land” place; a simple “reading” pattern connects the transition “fox appears to Pierson” with the transition after “seeking foxes”; a complete communication pattern, which separates emission from reception, surrounds the “fox taken by Pierson” place. See [25] for other examples.

The core action starts when the fox appears to Pierson (black box). Pierson acknowledges it, as he is seeking for foxes (and therefore is prone/motivated to see them), and commits himself to the catch, executing a specific sequence of actions. On the other side, Post was already trying to catch the fox. We modeled Post's institutional thinking (highlighted with the grey boxes), assumed to be similar to Livingston's interpretation. Post started the action as Pierson, but the Petri net shows also that he infers the institutional power to take it, before he attempts to (left box). The same mechanism is hereafter applied to Pierson's intervention (right box), but is *defeated* by a second mechanism which gives Post the exclusive right on the prey. From Post's perspective, the event of taking the fox is therefore not only a practical failure, but also a social failure.

6 Discussion

The paper introduces several elements towards an alternative framework for narratives:

- by integrating intentions and institutional concepts we are able to increase the deepness of the representation in order to model motivational and social aspects of stories, relevant to describe legal cases, but usually not explicitly targeted in other formal accounts;
- the use of Petri nets as underlying computational model makes the story-model execution a direct possibility, useful as well for real-time visual debugging and validating purposes. Other contributions referring to Petri net have been proposed for *story plot generation* in games [5] and for *narrative comprehension* [18, 12]. Although they have similarities with our approach (in particular the latter), the first describes only higher-level specifications, the second focuses with much further detail on narrative discourse components, while we focus on the acquisition of the “systematic” structural core of an interpreted story;
- *causations* and *expectations* are integrated in the same framework, overcoming mutual limitations of other frameworks (cf. [16]); furthermore, the connection with concepts as affordances, power, permission/obligation is, as far as we know, a novelty in the domain;
- the difference between expectation and actual outcome — which eventually defines *failure* or *success* for a character — is evaluated within the model itself. Therefore, the narrative analyst is not concerned anymore by this meta-interpretation, as occurs instead in *plot units* [13] or in *story intention graphs* (the “affectual” components) [9];
- rather than static *script*-like knowledge to be used in story-understanding, we focus on acquiring from modelers different interpretations of the fabula; assessing a case with alternative interpretations is just a daily practice in legal activity, and our representational framework aims to collect the correspondent mechanisms.

Obviously, the increase of knowledge requirements corresponds to increased effort in the modeling exercise. Ideally, this should not be a problem with end-users as law students and experts working on a legal scenario, because the visualization of the mechanisms produces also the direct effect of clarifying and validating the ideas of the modeler. In this direction, the choice of relying on visual programming practices aims to help the interaction with non-IT experts.

Moreover, we think there are two other reasons why the impact may be less critical than it seems. On the one hand, the elicitation of mechanisms is mostly targeting the *affordances* and *dispositions* related to actions/events, which are *the* fundamental components of practical rationality. We assume nobody should have particular problems with this part, if supported with an adequate acquisition platform. The composition of constraints, which is a more delicate and complex task, can be instead supported computationally (problems known in AI as *configuration* and *model-based diagnosis*). On the other hand, once a story is collected,

its mechanisms can be reused with another story. Evidently, in some cases part of those mechanisms would not directly apply; in such situation, the modeler will be obliged to define distinguishing features. This incremental, constructive approach reflects essentially the nature of *case law*. However, the overall impact of these potentially positive interactions with the knowledge acquisition bottleneck issue remains to be investigated in the future.

Some final remarks about a related domain. Traditionally, *legal case-based reasoning* binds the modeling of cases to *dimensions* (HYPO [3]) and *factors* (CATO [2]), i.e. concepts which translate legally significant aspects of the cases. The story behind the case and its construction is therefore neglected, apart from relevant components extracted by the analyst. The analyst is then the one responsible for finding the analogies and for placing the case in the right abstraction. Potentially, our framework could automatize part of this process, if the case-base is adequately rich and we implement an adequate measure of similarity.

References

- 1 Robert P. Abelson and Roger C. Schank. *Scripts, plans, goals and understanding*. Lawrence Erlbaum Associates, 1977.
- 2 Vincent Aleven. *Teaching Case-Based Argumentation Through a Model and Examples*. PhD thesis, University of Pittsburgh, 1997.
- 3 Kevin D. Ashley. Reasoning with cases and hypotheticals in HYPO. *International Journal of Man-Machine Studies*, 34(6):753–796, 1991.
- 4 Mieke Bal. *Narratology: Introduction to the Theory of Narrative*. University of Toronto Press, 2nd edition, 1997.
- 5 Daniel Balas, Cyril Brom, Adam Abonyi, and Jakub Gemrot. Hierarchical petri nets for story plots featuring virtual humans. *Proc. of 4th AI and Interactive Digital Entertainment Conf.*, pages 2–9, 2008.
- 6 Trevor J.M. Bench-Capon and Giovanni Sartor. A model of legal reasoning with cases incorporating theories and values. *Artificial Intelligence*, 150(1-2):97–143, 2003.
- 7 Floris Bex and Bart Verheij. Solving a Murder Case by Asking Critical Questions: An Approach to Fact-Finding in Terms of Argumentation and Story Schemes. *Argumentation*, 26(3):325–353, 2011.
- 8 Alexander Boer and Tom van Engers. An Agent-based Legal Knowledge Acquisition Methodology for Agile Public Administration. *ICAIL 2011: Proc. of the 13th Int. Conf. on Artificial Intelligence and Law*, June, 2011.
- 9 David K. Elson. DramaBank: Annotating Agency in Narrative Discourse. *LREC 2012: Proc. of the 8th Int. Conf. on Language Resources and Evaluation*, 2012.
- 10 Edward M. Forster. *Aspects of the novel*. Edward Arnold, London, 1927.
- 11 David Harel and Rami Marelly. *Come, let's play: Scenario-based programming using LSCs and the Play-Engine*, volume 1. Springer, 2003.
- 12 George Lakoff and Srinivas Narayanan. Toward a computational model of narrative. In *Computational Models of Narrative: Papers from AAAI Fall Symposium*, pages 21–28, 2010.
- 13 Wendy G. Lehnert. Plot units and narrative summarization. *Cognitive Science*, 5(4):293–331, 1981.
- 14 Charlotte Linde. Narrative and social tacit knowledge. *Journal of knowledge management*, 5(2):160–171, 2001.
- 15 Benedikt Löwe, E Pacuit, and S Saraf. Identifying the structure of a narrative via an agent-based logic of preferences and beliefs: Formalizations of episodes from CSI: Crime Scene Investigation. *MOCA 2009: Proc. of the 5th Int. Workshop on Modelling of Objects, Components and Agents*, 2009.

- 16 Benedikt Löwe. Methodological remarks about comparing formal frameworks for narratives. *Proc. of the 3rd Workshop in the Philosophy of Information*, pages 45–63, 2011.
- 17 Erik T. Mueller. Story understanding through multi-representation model construction. In *Text Meaning: Proc. of the HLT-NAACL 2003 Workshop*, pages 46–53, 2003.
- 18 Srinu Narayanan. *Knowledge-based Action Representations for Metaphor and Aspect (KARMA)*. PhD thesis, University of California at Berkeley, 1997.
- 19 Nancy Pennington and Reid Hastie. Reasoning in explanation-based decision making. *Cognition*, 49:123–163, 1993.
- 20 Vladimír Propp. *Morphology of the Folktale*. University of Texas Press, Austin, 1968.
- 21 David E. Rumelhart. Notes on a schema for stories. In D. G. Bobrow & A. Collins, editor, *Representation and understanding: Studies in cognitive science*. 1975.
- 22 Searle, J.R., Vanderveken, D.: Introduction to the Theory of Speech Acts. In: Foundations of Illocutionary Logic, vol. 1. Cambridge University Press (1985)
- 23 Giovanni Sileno, Alexander Boer, and Tom van Engers. Towards a Representational Model of Social Affordances from an Institutional Perspective. In *SOCIAL.PATH 2014: Proc. of the 2nd Workshop on Computational Social Science and Social Computer Science: Two Sides of the Same Coin*, 2014.
- 24 Giovanni Sileno, Alexander Boer, and Tom van Engers. Towards a computational model for institutional scenarios. In *BNAIC 2013: Proc. of the 25th Benelux Conf. Artificial Intelligence*, pages 183–190, 2013.
- 25 Giovanni Sileno, Alexander Boer, and Tom van Engers. From Inter-Agent to Intra-Agent Representations: Mapping Social Scenarios to Agent-Role Descriptions. In *ICAART 2014: Proc. of the 6th Int. Conf. Agents and Artificial Intelligence*, 2014.
- 26 Giovanni Sileno, Alexander Boer, and Tom van Engers. Implementing Explanation-Based Argumentation using Answer Set Programming. In *ArgMAS 2014: Proc. of the 11th Int. Workshop on Argumentation in Multi-Agent Systems*, 2014.
- 27 Katharine Young. Ontological puzzles about narrative. *Poetics*, 13, 1984.