

Computability of Narrative

Loizos Michael

Open University of Cyprus
loizos@ouc.ac.cy

Abstract

Among the many aspects of human intelligence that currently elude the simulation by machines is that of story understanding. Although many theories of narrative have been proposed, several processes pertaining to narrative remain inadequately formalized and, hence, beyond full mechanization. This work proposes a general formal framework that attempts to make precise such processes and related notions, with first and foremost that of what constitutes a narrative. Emphasis is placed on identifying certain premises that narratives are expected to adhere to, and deriving the formal implications that these have in terms of the computability of the various relevant notions. Among others, it is established that checking whether a discourse is a narrative is decidable, and that narratives can be computably enumerated and, hence, unambiguously indexed.

Prologue and Sneak Preview

Narratives are ubiquitous in everyday life, and play an important role in the process of communication between humans. Indeed, communication often involves more than simply the transmission of a piece of information, and places emphasis on the transmission of a narrative, a sequence of statements that together aim to “tell a story”. This simple observation already brings about a multitude of questions regarding the characterization, generation, and manipulation of narratives. Providing answers to these questions is directly relevant to two long-term scientific endeavors: (i) understanding aspects of human cognition and social interaction; (ii) developing machines that interact with humans in a natural manner. We propose an abstract framework that makes progress in answering some of these questions.

This work proposes a computational framework within which narratives can be formally studied and understood. Alongside the particular framework, this work also identifies and makes explicit certain premises that we contend are central, and perhaps necessary, in any attempt to understand narratives. These premises are motivated through a discussion in the following section, and are formalized as part of the proposed framework later on. The first premise is that narratives cannot be understood out of context. What counts as a narrative in one environment does not necessarily count as such in another. This, then, necessitates a means to encode and reason with knowledge about environments. The second premise is that this knowledge need not be factual, but may

capture beliefs or commonsense knowledge that an agent possesses about its environment. It, then, becomes clear that representing and reasoning with commonsense knowledge is, in fact, a prerequisite for understanding narratives. Indeed, our approach to understanding narratives in a given environment is to view them as projections of models of an agent’s commonsense theory of that environment. Recognizing or generating narratives reduces, then, to computing and manipulating models of a commonsense theory.

Emphasis in this work is placed on investigating the computability of certain aspects of narrative. Is there a definition for narrative? Is checking narratives according to this definition decidable? Are there more relaxed notions of narrative that are still acceptable? Do such less stringent forms of narrative always exist? Can we identify among certain narratives which one is more preferred? Can narratives be indexed so that their storage and retrieval is possible? Does there exist a computable set containing all possible narratives? We next answer these questions in the affirmative.

What Counts as a Narrative?

Perhaps the most important and immediate aspect of a narrative is that of being a discourse, or a sequence of statements: “statements” as in making a claim that something holds or occurs, and “sequence” as in imposing an ordering over these statements. Beyond that, the necessary conditions of what constitutes a narrative get a bit more difficult to pin down. For instance, consider the following discourse:

The judge finds John Smith guilty of a first-degree murder of his wife. Mr. Smith is sentenced for his crime, and spends 25 years in a prison. John Smith has an argument with his wife, and stabs her with a knife.

Would this discourse count as a narrative? Presumably, most people would be confident in saying that it would not, as it violates temporal continuity. The last statement should have been first! Although this seems like a reasonable response, a closer look is warranted. Let us assume for a moment, the existence of a planet in a galaxy far far away, where technology is so advanced that it is possible to determine with certainty from one’s genes the crimes that one will make. Fittingly, the planet’s judicial system is appropriately adapted to this technological ability. According to this judicial system, every resident is prosecuted once she turns 23 years old,

on account of all the crimes that she has committed or will commit for the duration of her lifetime. The resident is then imprisoned for an appropriate period of time, after which she is free to lead the rest of her life, including committing any crimes (for which she has already paid), without being prosecuted again. Fictional as this planet may be, we need to ask ourselves whether the discourse above would have counted as a narrative, had we been residents of this planet. We argue that the answer to this hypothetical question is affirmative!

The considered thought experiment suggests that whether a discourse counts as a narrative cannot be decided unconditionally and outside some given context. Instead, whether a discourse counts as a narrative can be decided only *with respect to a domain*. We use the term “domain” to mean a set of constraints that are expected to be obeyed in the world to which the domain corresponds. This set could comprise, for instance, causal laws that encode how properties evolve over time, static laws that encode how properties relate to each other at each point in time, and even specific facts or events that are known to hold or occur in the corresponding world. With respect to such a domain, then, a discourse is a narrative if it does not violate any of the domain constraints.

Let us now revisit the discourse above, and consider once more whether it could count as a narrative with respect to our own world. The response given earlier was that this is not the case. But, are there *any* circumstances under which we would be willing to accept the opposite response? We argue that, in fact, there are such circumstances. If Mr. Smith’s wife mentioned in the third statement is not the same person as Mr. Smith’s wife mentioned in the first statement, then no constraints are violated, and, hence, the discourse should count as a narrative. Why are we, then, more reluctant to accept this discourse as a narrative with respect to our world than with respect to the fictional world discussed earlier? Why was our initial response that this discourse is *not* a narrative with respect to our world, and why did we presumably feel so confident in offering that initial response, even though it is *possible* for such a discourse to accurately represent some actual sequence of events in the real world?

We suggest that an answer comes from the following realization: a domain with respect to which a narrative is interpreted, contains not only factual knowledge, but also beliefs or assumptions that an agent has or makes with respect to its environment. One such belief with respect to our world could be, for instance, that each person has only one spouse during their lifetime. Along with factual knowledge, such beliefs and assumptions act as constraints of the domain also, and as such they are expected to be respected by narratives. In cultures where humans routinely have multiple spouses during their lifetime, for instance, the domain would not have contained the corresponding belief-derived constraint, and Mr. Smith could be killing one wife after another within the same discourse without this prompting us to feel reluctant to count it as a narrative. In this sense, then, we suggest that the notion of a “narrative” may differ significantly across cultures, at least on those aspects that are tied to social norms followed by different cultures. On the other hand, aspects of narratives that relate to physical properties of our world, and are, hence, constrained in the same manner

in domains across all cultures, are expected to be the same.

Beyond the cultural-dependent nature of narratives, a second aspect of narrative also becomes apparent from the discussion above. In cases where certain beliefs are violated by a discourse, we might still be willing to accept the discourse as a narrative, even though somewhat reluctantly. This, we believe, is a result of recognizing, as humans, that beliefs are simply that: beliefs. Thus, they are not indisputable, and the constraints they give rise to hold only by default. If, by retracting some of these default assumptions, a discourse no longer violates the (remaining) constraints of a domain, then we may accept it as a narrative. We suggest that the more default assumptions within a domain need to be retracted in order for the discourse to respect the (remaining) constraints of a domain, the more reluctant we are in counting the discourse as a narrative. This degree of reluctance can be seen to determine a degree of “narrativeness” of a discourse. It is our view, therefore, that the notion of a discourse “being a narrative” should not be thought of as a boolean property, but rather as a graded property with a degree of certainty.

A Computational Framework

Following our discussion in the preceding section, we formally define next what is a discourse, what is a domain, and when the former is a narrative with respect to the latter. In the context of this work we employ a simplified syntax and semantics for the framework that we develop, as this is already sufficient to illustrate the main aspects of narratives, and to formally represent a number of narratives. Richer syntax and semantics can be substituted for what we employ below, without affecting the main premises that we consider.

We assume the existence of a propositional language $\langle \mathcal{F}, \mathcal{A} \rangle$, where \mathcal{F} is a finite set of fluents that name basic properties of the environment, and \mathcal{A} is a finite set of actions that name actions or events in the environment. Implicit in language $\langle \mathcal{F}, \mathcal{A} \rangle$ is the standard set of logical connectives and the entailment operator \models of Propositional Calculus. In order to keep the presentation simple, we assume throughout that representations employ language $\langle \mathcal{F}, \mathcal{A} \rangle$. References to a time-line $\langle \mathcal{T}, \preceq^t \rangle$ assume the existence of a fixed countable set \mathcal{T} of time-points, and a total ordering \preceq^t over \mathcal{T} . For concreteness, we shall henceforth take \mathcal{T} to be the set of non-negative integers, and \preceq^t to be the standard ordering over integers. In particular, we shall use the notation $T+1$ to refer to the unique time-point following time-point $T \in \mathcal{T}$.

Our first definition formalizes the notion of a “discourse” as a partially ordered set of events and facts.

Definition 1 (Discourse). A *discourse* is a triple $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ comprising a finite set \mathcal{C} of clauses of the form $\text{occurs}(A, S)$ and $\text{holds}(L, S)$, a finite set \mathcal{S} of states, and an acyclic partial ordering \preceq^s over \mathcal{S} , where $A \in \mathcal{A}$, L is a literal over \mathcal{F} , and $S \in \mathcal{S}$. A discourse is *singular* if it contains $\text{holds}(F, S)$ and $\text{holds}(\neg F, S)$ for some fluent F and state S .

A discourse need not necessarily refer to absolute time-points. It is, therefore, possible in the general case to embed a discourse in a time-line in a number of ways, as long as the partial order of events and facts in the discourse is respected.

Definition 2 (Embedding). An *embedding* of a discourse $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ in a time-line $\langle \mathcal{T}, \preceq^t \rangle$ is a set of clauses of the form $\text{occurs}(A, T)$ and $\text{holds}(L, T)$ that results by substituting a time-point in \mathcal{T} for each state in \mathcal{S} , so that if the time-points $T_1, T_2 \in \mathcal{T}$ are substituted, respectively, for states $S_1, S_2 \in \mathcal{S}$ such that $S_1 \preceq^s S_2$, then $T_1 \preceq^t T_2$.

To represent and reason about the context with respect to which a discourse is to be interpreted, we introduce the notion of a “domain”. The following two definitions capture the syntax and semantics of domains, and closely follow typical definitions for domains found in the literature of commonsense reasoning about actions and change.

We suggest that sufficient expressivity can be achieved by allowing domains to encode four different types of knowledge: action occurrences, observation of facts, static constraints, and causal laws (or temporal constraints).

Definition 3 (Domain). A *domain* over a time-line $\langle \mathcal{T}, \preceq^t \rangle$ is a finite set \mathcal{D} of clauses of the form $\text{occurs}(A, T)$, $\text{holds}(L, T)$, $\text{static}(\phi)$, $\text{causes}(\psi, L)$, where $A \in \mathcal{A}$, L is a literal over \mathcal{F} , ϕ is a formula over \mathcal{F} , ψ is a formula over $\mathcal{A} \cup \mathcal{F}$, and $T \in \mathcal{T}$.

A semantics to the different pieces of knowledge is given in model-theoretic terms, by insisting that any sequence of states that represent the evolution of an environment should be such that: (i) actions are executed at the time-point of their occurrence; (ii) facts hold at the time-point of their observation, and static constraints are satisfied at every time-point; (iii) changes that are caused at some time-point are brought about¹ immediately afterwards; and (iv) properties of the environment that are not caused to change persist.

Definition 4 (Model). Consider a domain \mathcal{D} over a time-line $\langle \mathcal{T}, \preceq^t \rangle$. An *assignment* to $\langle \mathcal{T}, \preceq^t \rangle$ is a mapping M of each pair of $X \in \mathcal{A} \cup \mathcal{F}$ and $T \in \mathcal{T}$ to a truth-value $M(X, T)$. The truth-assignment over $\mathcal{A} \cup \mathcal{F}$ that is induced by projecting / restricting M to a given time-point $T \in \mathcal{T}$ is denoted by $M(T)$. A *model* of \mathcal{D} is an assignment M to $\langle \mathcal{T}, \preceq^t \rangle$ such that for each $A \in \mathcal{A}$, each literal L over \mathcal{F} , each formula ϕ over \mathcal{F} , and each $T \in \mathcal{T}$, the following conditions hold:

- (i) $M(T) \models A$ if and only if $\text{occurs}(A, T) \in \mathcal{D}$.
- (ii) $M(T) \models L$ if $\text{holds}(L, T) \in \mathcal{D}$.
 $M(T) \models \phi$ if $\text{static}(\phi) \in \mathcal{D}$.
- (iii) $M(T+1) \models L$ if $M(T) \models \psi$ for some ψ such that $\text{causes}(\psi, L) \in \mathcal{D}$.
- (iv) $M(T+1) \models L$ if $M(T) \models L$ and $M(T) \not\models \psi$ for every ψ such that $\text{causes}(\psi, \neg L) \in \mathcal{D}$.

A domain \mathcal{D} is *consistent* if there exists a model of \mathcal{D} .

Using the notion of a “domain” to represent the context of a discourse, and hence the set of constraints that are expected to be satisfied, a narrative is defined to be a discourse that is consistent with this set of constraints.

¹ Although domains contain only deterministic causal laws, default domains, which shall be introduced later, are sufficiently expressive to capture non-deterministic causal laws as well. We shall not discuss the issue of determinism in causal laws further, as the emphasis of this work is not on the semantics of domains per se.

Definition 5 (Narrative). Consider a domain \mathcal{D} over a time-line $\langle \mathcal{T}, \preceq^t \rangle$. A *narrative* w.r.t. \mathcal{D} is a discourse $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ such that there exists an embedding of $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ in $\langle \mathcal{T}, \preceq^t \rangle$, whose union with \mathcal{D} results in a consistent domain.

Note that if the time-line $\langle \mathcal{T}, \preceq^t \rangle$ is infinite, then there are infinitely many embeddings of a discourse $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ in $\langle \mathcal{T}, \preceq^t \rangle$. This immediately raises the question of whether being a narrative with respect to a given domain \mathcal{D} is even a decidable property! Although answering this question is certainly not trivial, the following result shows that the answer is affirmative, and that decidability of narrative is ultimately a result of the finiteness of the set \mathcal{F} of fluents, the set \mathcal{C} of clauses of the discourse $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$, and the domain \mathcal{D} .

Theorem 1 (Decidability of Narrative). Consider a domain \mathcal{D} over a time-line $\langle \mathcal{T}, \preceq^t \rangle$, and a discourse $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$. Then, checking whether $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ is a narrative w.r.t. \mathcal{D} is decidable.

Proof. Let T_0 denote the largest time-point referenced in the clauses of domain \mathcal{D} . For each embedding \mathcal{E} of $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ in $\langle \mathcal{T}, \preceq^t \rangle$, let $T_{\mathcal{E}}$ denote the set of time-points that are referenced in the clauses of \mathcal{E} and also follow time-point T_0 .

We shall first establish the truth of the following lemma: if there exists an embedding \mathcal{E}_1 of $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ in $\langle \mathcal{T}, \preceq^t \rangle$ such that $\mathcal{D} \cup \mathcal{E}_1$ is consistent and $\max(T_{\mathcal{E}_1}) \geq T_0 + (2^{|\mathcal{F}|} + 2) \cdot |\mathcal{C}|$, then there exists an embedding \mathcal{E}_2 of $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ in $\langle \mathcal{T}, \preceq^t \rangle$ such that $\mathcal{D} \cup \mathcal{E}_2$ is consistent and $\max(T_{\mathcal{E}_2}) < \max(T_{\mathcal{E}_1})$.

Consider an embedding \mathcal{E}_1 as stated above in the lemma. Since $\mathcal{D} \cup \mathcal{E}_1$ is consistent, let M_1 be a model of $\mathcal{D} \cup \mathcal{E}_1$.

Note that the set $\{T_0\} \cup T_{\mathcal{E}_1}$ contains at most $|\mathcal{C}| + 1$ time-points, spanning an interval of size at least $(2^{|\mathcal{F}|} + 2) \cdot |\mathcal{C}| + 1$. Since T_0 is at the beginning of this interval, there exist time-points $T_1, T_2 \in \{T_0\} \cup T_{\mathcal{E}_1}$ such that $T_2 - T_1 \geq 2^{|\mathcal{F}|} + 2$, and no time-point $T \in \{T_0\} \cup T_{\mathcal{E}_1}$ is such that $T_1 < T < T_2$.

Since no time-point T between T_1 and T_2 is referenced in the clauses of $\mathcal{D} \cup \mathcal{E}_1$, the projection $M_1(T)$ takes one of at most $2^{|\mathcal{F}|}$ different values. Since there are at least $2^{|\mathcal{F}|} + 1$ time-points between T_1 and T_2 , there exist time-points $T'_1, T'_2 : T_1 < T'_1 < T'_2 < T_2$ such that $M_1(T'_1) = M_1(T'_2)$.

Consider the embedding \mathcal{E}_2 of $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ in $\langle \mathcal{T}, \preceq^t \rangle$ obtained from \mathcal{E}_1 by substituting $T - (T'_2 - T'_1)$ for every time-point $T > T'_2$ that is referenced in the clauses of \mathcal{E}_1 . It follows that $\max(T_{\mathcal{E}_2}) = \max(T_{\mathcal{E}_1}) - (T'_2 - T'_1) < \max(T_{\mathcal{E}_1})$. Consider, also, an assignment M_2 to $\langle \mathcal{T}, \preceq^t \rangle$ defined such that $M_2(T) = M_1(T)$ for every time-point $T \leq T'_1$, and $M_2(T) = M_1(T + (T'_2 - T'_1))$ for every time-point $T > T'_1$. By Definition 4, M_2 is a model of $\mathcal{D} \cup \mathcal{E}_2$, and, hence, $\mathcal{D} \cup \mathcal{E}_2$ is consistent. From the above, then, the lemma holds.

It now follows that when checking whether $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ is a narrative w.r.t. \mathcal{D} , it suffices to check whether there exists an embedding of $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ in $\langle \mathcal{T}, \preceq^t \rangle$ that does not reference time-points from $T_0 + (2^{|\mathcal{F}|} + 2) \cdot |\mathcal{C}|$ onwards, and whose union with \mathcal{D} is consistent. If such an embedding exists, then $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ is a narrative w.r.t. \mathcal{D} . Otherwise, by repeated application of the lemma above, there exists no embedding of $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ in $\langle \mathcal{T}, \preceq^t \rangle$ whose union with \mathcal{D} is consistent, and, therefore, $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ is not a narrative w.r.t. \mathcal{D} .

To complete the proof, then, it remains to establish that checking the entire set of time-restricted embeddings mentioned in the preceding paragraph can be done in finite time.

Since the set \mathcal{C} of clauses of the discourse $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ is finite, there are finitely many states that are referenced in the clauses of \mathcal{C} . Since the time-points preceding $T_0 + (2^{|\mathcal{F}|} + 2) \cdot |\mathcal{C}|$ are also finite, there are finitely many time-restricted embeddings of $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ in $\langle \mathcal{T}, \preceq^t \rangle$ to be considered. Each of these embeddings \mathcal{E} is considered in turn, and a new domain $\mathcal{D} \cup \mathcal{E}$ is computed, which must be checked for consistency. Since both \mathcal{D} and \mathcal{E} are finite, so is $\mathcal{D} \cup \mathcal{E}$. By an argument analogous to the one in the proof of the lemma above, it can be shown that some projection of any model of $\mathcal{D} \cup \mathcal{E}$ will eventually be repeated. Hence, checking that the conditions of Definition 4 are satisfied can be restricted to time-points T up to $T_0 + (2^{|\mathcal{F}|} + 2) \cdot |\mathcal{C}| + 2^{|\mathcal{F}|}$. Therefore, consistency of $\mathcal{D} \cup \mathcal{E}$ can be checked in finite time. The claim follows. \square

Overall, our framework suggests the following approach for recognizing whether a given discourse is a narrative: Identify the context with respect to which the discourse is to be understood and encode it as a domain. Map the discourse into a set of events and facts, and embed those (in only finitely many ways) in the domain representing the context. If the resulting extended domain is consistent for one of the embeddings, then the given discourse can exist without violating the constraints of its context. This discourse is, then (and only then), a narrative with respect to that context.

Absolute versus Preferred Narratives

Definition 5 is absolute in insisting that all the domain constraints are to be satisfied. We have, however, argued earlier that “narrativeness” is not necessarily a boolean property, but that different discourses may have a different degree of “narrativeness” depending on how many of the domain constraints they satisfy. We next formalize this view of things.

Recall from the preceding section that certain constraints need not necessarily be satisfied by narratives on the grounds that they represent beliefs, and not factual knowledge. The following definition aims to make the distinction between strict and defeasible knowledge precise. In the interest of generality, our framework does not determine the details of how such a distinction is made. Instead, it abstractly captures the requirement that all strict knowledge be satisfied by letting a default domain determine the set Δ of subsets of \mathcal{D} that correspond to such strict knowledge. At the same time, the requirement that defeasible knowledge be satisfied to the extent possible is captured by letting a default domain impose a preference \preceq^d over Δ .

Definition 6 (Default Domain). A *default domain over a time-line* $\langle \mathcal{T}, \preceq^t \rangle$ is a triple $\langle \mathcal{D}, \Delta, \preceq^d \rangle$ comprising a domain \mathcal{D} over $\langle \mathcal{T}, \preceq^t \rangle$, a subset $\Delta \subseteq 2^{\mathcal{D}}$ of domains over $\langle \mathcal{T}, \preceq^t \rangle$, and a transitive preference relation \preceq^d over Δ .

In typical domains, \preceq^d may be equated to the set-theoretic inclusion relation, so that satisfying more (in a set-theoretic sense) pieces of defeasible knowledge is always preferred. Of course, much more generality is possible. The relation may be extended so that subsets that are incomparable with

respect to the set-theoretic inclusion may have a preference imposed among them depending on the types of defeasible knowledge that each subset contains. In certain domains, for instance, one may prefer that defeasible static knowledge be satisfied at the expense of defeasible causal laws. At a finer grade, one may prefer that a particular piece of static knowledge be satisfied at the expense of another. In certain situations, such preferences may even violate set-theoretic inclusion. Ultimately, the definition of \preceq^d is domain-dependent.

The preference over which sub-domains (effectively, subsets of defeasible knowledge) of a domain are to be satisfied, imposes a natural preference over which discourses are to be considered narratives with respect to the domain. First, only discourses that satisfy *some* sub-domain are considered. These candidate narratives essentially correspond to those discourses that satisfy at least the strict knowledge of the domain. Among all those candidate narratives, preference is given to those that satisfy more preferred sub-domains (and, hence, subsets of defeasible knowledge).²

Definition 7 (Preferred Narrative). Consider a default domain $\langle \mathcal{D}, \Delta, \preceq^d \rangle$ over a time-line $\langle \mathcal{T}, \preceq^t \rangle$. A discourse $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ is a *candidate narrative w.r.t.* $\langle \mathcal{D}, \Delta, \preceq^d \rangle$ if it is a narrative w.r.t. some domain that belongs in Δ .

Consider a set \mathcal{N} of candidate narratives w.r.t. $\langle \mathcal{D}, \Delta, \preceq^d \rangle$. A discourse $\langle \mathcal{C}_1, \mathcal{S}_1, \preceq_1^s \rangle$ is a *preferred narrative of \mathcal{N} w.r.t.* $\langle \mathcal{D}, \Delta, \preceq^d \rangle$ if the following conditions hold:

- (i) $\langle \mathcal{C}_1, \mathcal{S}_1, \preceq_1^s \rangle \in \mathcal{N}$ is a narrative w.r.t. a domain $\mathcal{D}_1 \in \Delta$.
- (ii) for every discourse $\langle \mathcal{C}_2, \mathcal{S}_2, \preceq_2^s \rangle \in \mathcal{N}$ that is a narrative w.r.t. a domain $\mathcal{D}_2 \in \Delta$, if $\mathcal{D}_1 \preceq^d \mathcal{D}_2$ then $\mathcal{D}_2 \preceq^d \mathcal{D}_1$.

The following result establishes that a preferred narrative always exists. In simple terms, this implies that given any non-empty set of discourses that are guaranteed to satisfy at least the strict knowledge of a domain, one may always identify (at least) one of those discourses that is most easily accepted as being a narrative with respect to the domain. Note that such a preferred narrative need not necessarily satisfy all pieces of defeasible knowledge of the domain.

Theorem 2 (Existence of Preferred Narrative). Consider a default domain $\langle \mathcal{D}, \Delta, \preceq^d \rangle$ over a time-line $\langle \mathcal{T}, \preceq^t \rangle$, and a non-empty set \mathcal{N} of candidate narratives w.r.t. $\langle \mathcal{D}, \Delta, \preceq^d \rangle$. Then, there exists a preferred narrative of \mathcal{N} w.r.t. $\langle \mathcal{D}, \Delta, \preceq^d \rangle$.

Proof. Since \mathcal{N} is not empty, consider any discourse $\langle \mathcal{C}_1, \mathcal{S}_1, \preceq_1^s \rangle \in \mathcal{N}$ that is a narrative w.r.t. a domain $\mathcal{D}_1 \in \Delta$.

²Discourses are compared w.r.t. a single *common* default domain, since the preferences based on which the comparison is made are inherently domain-specific. If, however, discourses need to be compared w.r.t. multiple default domains, then one of two possible natural approaches can be followed: (i) The default domains can be first merged into one, the preference relation of which can be then appropriately extended to include explicit preferences across sub-domains of the original default domains. (ii) A multi-dimensional comparison of discourses can be made, with each default domain w.r.t. which discourses are to be compared contributing to one of the dimensions. Preference can be then given to those discourses that belong to the Pareto frontier of this multidimensional space.

- 1: Construct a directed graph $G = \langle V, E \rangle$ with vertex set $V = \Delta$, and edge set $E = \{(\mathcal{D}_2, \mathcal{D}_1) \mid \mathcal{D}_1 \preceq^d \mathcal{D}_2 \text{ and } \mathcal{D}_2 \not\preceq^d \mathcal{D}_1\}$.
- 2: Topologically sort the vertex set V according to G , and let V_1, V_2, \dots, V_n be the resulting partitioning of V , with lower-indexed sets containing vertices that topologically precede those in higher-indexed sets.
- 3: For each integer value of j from 1 to n , set $\mathcal{N}_j := \emptyset$.
- 4: For each discourse $\langle \mathcal{C}_i, \mathcal{S}_i, \preceq_i^s \rangle \in \mathcal{N}$, do:
- 5: For each integer value of j from 1 to n , do:
- 6: For each domain $\mathcal{D}_k \in V_j$, do:
- 7: If $\langle \mathcal{C}_i, \mathcal{S}_i, \preceq_i^s \rangle$ is a narrative w.r.t. \mathcal{D}_k , then do:
- 8: Set $\mathcal{N}_j := \mathcal{N}_j \cup \{\langle \mathcal{C}_i, \mathcal{S}_i, \preceq_i^s \rangle\}$.
- 9: Break and continue at Step 4.
- 10: Return $\mathcal{N}_1, \mathcal{N}_2, \dots, \mathcal{N}_n$, and terminate.

Algorithm 1. An algorithm that sorts a given set of candidate narratives according to the preference relation determined by the corresponding default domain.

If it is not a preferred narrative of \mathcal{N} w.r.t. $\langle \mathcal{D}, \Delta, \preceq^d \rangle$, then by Definition 7 there exists a discourse $\langle \mathcal{C}_2, \mathcal{S}_2, \preceq_2^s \rangle \in \mathcal{N}$ that is a narrative w.r.t. a domain $\mathcal{D}_2 \in \Delta$, and such that $\mathcal{D}_1 \preceq^d \mathcal{D}_2$ and $\mathcal{D}_2 \not\preceq^d \mathcal{D}_1$. Consider this discourse and repeat the process. By the transitivity of \preceq^d , no domain will be encountered twice in this process, and by the finiteness of \mathcal{D} , and hence Δ , the process will eventually stop. The last discourse considered will be a preferred narrative. \square

The existence of a preferred narrative in every non-empty set of candidate narratives, implies the ability to order candidate narratives according to how preferred they are as narratives. More precisely, a set \mathcal{N} of candidate narratives w.r.t. a default domain $\langle \mathcal{D}, \Delta, \preceq^d \rangle$ can be partitioned into subsets $\mathcal{N}_1, \mathcal{N}_2, \dots, \mathcal{N}_n$ such that: the least-indexed non-empty subset \mathcal{N}_{j_1} contains all the preferred narratives of \mathcal{N} w.r.t. $\langle \mathcal{D}, \Delta, \preceq^d \rangle$; when all discourses in \mathcal{N}_{j_1} are removed from \mathcal{N} , the next least-indexed non-empty subset \mathcal{N}_{j_2} contains all the preferred narratives of the resulting \mathcal{N} w.r.t. $\langle \mathcal{D}, \Delta, \preceq^d \rangle$; and so on for the remaining subsets. Algorithm 1 illustrates how such a partition can be constructed, and Theorem 3 establishes that the particular algorithm works as expected.

Theorem 3 (Narrative Sorting in Preferred Order). *Given a default domain $\langle \mathcal{D}, \Delta, \preceq^d \rangle$, and a set \mathcal{N} of candidate narratives w.r.t. $\langle \mathcal{D}, \Delta, \preceq^d \rangle$, Algorithm 1 terminates and returns a partition $\mathcal{N}_1, \mathcal{N}_2, \dots, \mathcal{N}_n$ of \mathcal{N} s.t. each discourse in \mathcal{N}_j is a preferred narrative of $\mathcal{N} \setminus \bigcup_{t=0}^{j-1} \mathcal{N}_t$ w.r.t. $\langle \mathcal{D}, \Delta, \preceq^d \rangle$.*

Proof. Consider the algorithm. First recall that Step 7 is decidable by Theorem 1. Second observe that each discourse in \mathcal{N} is considered and checked for being a narrative with respect to every domain in Δ . Since each discourse in \mathcal{N} is, in fact, a candidate narrative w.r.t. $\langle \mathcal{D}, \Delta, \preceq^d \rangle$, it follows that the condition at Step 7 will succeed exactly once for

each discourse, and hence the algorithm will return a partition $\mathcal{N}_1, \mathcal{N}_2, \dots, \mathcal{N}_n$ of \mathcal{N} and will terminate as needed. It remains to show that the partition has the claimed property.

Assume, by way of contradiction, that there exists some discourse $\langle \mathcal{C}_1, \mathcal{S}_1, \preceq_1^s \rangle \in \mathcal{N}_{j_1}$ that is not a preferred narrative of $\mathcal{N} \setminus \bigcup_{t=0}^{j_1-1} \mathcal{N}_t$ w.r.t. $\langle \mathcal{D}, \Delta, \preceq^d \rangle$. Then, $\langle \mathcal{C}_1, \mathcal{S}_1, \preceq_1^s \rangle$ is a narrative w.r.t. a domain $\mathcal{D}_{k_1} \in V_{j_1}$. By Definition 7, there exists a discourse $\langle \mathcal{C}_2, \mathcal{S}_2, \preceq_2^s \rangle \in \mathcal{N} \setminus \bigcup_{t=0}^{j_1-1} \mathcal{N}_t$ that is a narrative w.r.t. a domain $\mathcal{D}_{k_2} \in \Delta$, such that $\mathcal{D}_{k_1} \preceq^d \mathcal{D}_{k_2}$ and $\mathcal{D}_{k_2} \not\preceq^d \mathcal{D}_{k_1}$. Let j_2 be such that $\mathcal{D}_{k_2} \in V_{j_2}$. Since $\langle \mathcal{C}_2, \mathcal{S}_2, \preceq_2^s \rangle \in \mathcal{N} \setminus \bigcup_{t=0}^{j_1-1} \mathcal{N}_t$, $\mathcal{D}_{k_2} \in V \setminus \bigcup_{t=0}^{j_1-1} V_t$; thus, $j_2 \geq j_1$. Since $\mathcal{D}_{k_1} \preceq^d \mathcal{D}_{k_2}$ and $\mathcal{D}_{k_2} \not\preceq^d \mathcal{D}_{k_1}$, $(\mathcal{D}_{k_2}, \mathcal{D}_{k_1}) \in E$; thus, $j_2 < j_1$. A contradiction, as needed. \square

Thus, given a set of discourses and a default domain acting as the context within which these discourses are to be understood, our framework suggests the following approach for determining which of the discourses are acceptable as narratives, and to what extent they are so: Discard all discourses that are not candidate narratives; these are not acceptable as narratives, since they contradict the strict constraints of the domain. Enumerate the remaining discourses according to their preferred order, and use this order to determine the extent of acceptability of each discourse as narrative. Note, for instance, that \mathcal{N}_1 may be empty, indicating that the most preferred over all discourses in \mathcal{N} is not as acceptable as a narrative as some other discourse (not in \mathcal{N}) could be. Thus, the returned order is not indicative only of the relative preference of the given discourses, but also of a more absolute and global notion of preference. Hence, one could use the index of \mathcal{N}_j within which each given discourse appears as an absolute (for the particular context) indicator of acceptability, with 1 indicating that a discourse is maximally acceptable as a narrative, and some integer n (which might be as large as $2^{|\mathcal{D}|}$) indicating that a discourse is minimally acceptable as a narrative (but acceptable nonetheless, since it satisfies at least the strict constraints of the domain).

Generating and Indexing Narratives

We have focused, so far, on the problem of recognizing narratives given access to discourses. Equally interesting is the problem of generating narratives without access to any discourses. As in the case of recognition, context is important.

Consider, first, the problem of generating narratives w.r.t. a consistent domain \mathcal{D} . Recall that a narrative is expected to satisfy all the constraints of \mathcal{D} . Recall, also, that models of \mathcal{D} are truth-assignments that satisfy all the clauses of \mathcal{D} across a time-line. This immediately suggests the following: Construct a model M of \mathcal{D} , and for some number of choices of $X \in \mathcal{A} \cup \mathcal{F}$ and $T \in \mathcal{T}$, translate the induced truth-values $M(X, T)$ into a set \mathcal{C} of corresponding clauses of the form $\text{occurs}(X, T)$, $\text{holds}(X, T)$, $\text{holds}(\neg X, T)$. Replace time-points with states giving rise to a set \mathcal{S} of states, and impose an ordering \preceq^s over the states so that it is consistent with the ordering of the time-points. This gives the discourse $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$, which, by construction, is a narrative w.r.t. \mathcal{D} .

Although this simple process can be used to generate narratives, the more interesting question is whether there is a

way to generate *all* narratives. To answer the question, note first that even though \mathcal{F} and \mathcal{A} are finite, there are infinitely many discourses, and hence narratives (in the same way that there are infinitely many possible novels despite there being only a finite number of letters). Thus, the best we can hope for is to show that the set of all narratives can be enumerated.

Many enumeration orderings of narratives $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ are possible, but we suggest the following natural ordering as a canonical one. First, enumerate narratives in order of increasing $|\mathcal{S}|$. Note that the set \mathcal{S} of states can be taken to be uniquely determined by its size, since renaming the states does not affect the properties of a discourse or a narrative. Therefore, without loss of generality, we may assume that for each narrative, $\mathcal{S} = \{S_1, S_2, \dots, S_{|\mathcal{S}|}\}$. $|\mathcal{S}|$ can, then, be taken to encode the temporal length of the narrative.

For narratives of the same temporal length, the order can be determined by considering \preceq^s . Note that \preceq^s is effectively a binary matrix of size $|\mathcal{S}| \times |\mathcal{S}|$, with the binary value of entry (i, j) indicating whether $S_i \preceq^s S_j$ holds or not. Such a matrix has an immediate representation as a non-negative integer (e.g., by reading its entries top to bottom and left to right, as a single binary number). This integer can, then, be taken to encode the temporal structure of the narrative.

Among narratives of the same temporal length, and the same temporal structure, the order can be determined by considering the actions that occur. Since \mathcal{A} is finite, one can impose some order on $\mathcal{A} = \{A_1, A_2, \dots, A_{|\mathcal{A}|}\}$ (e.g., lexicographic), and use a binary vector of size $|\mathcal{A}|$ to determine which actions occur at some specific state. The action occurrences across all states in \mathcal{S} can be encoded as a binary matrix of size $|\mathcal{A}| \times |\mathcal{S}|$, with the binary value of entry (i, j) indicating whether $\text{occurs}(A_i, S_j) \in \mathcal{C}$ holds or not. Such a matrix has an immediate representation as a non-negative integer (e.g., by reading its entries top to bottom and left to right, as a single binary number). This integer can, then, be taken to encode the active plot of the narrative.

Finally, for those narratives that share all the characteristics mentioned above, the order can be determined by considering the facts that are observed. Since \mathcal{F} is finite, one can impose some order on $\mathcal{F} = \{F_1, F_2, \dots, F_{|\mathcal{F}|}\}$ (e.g., lexicographic), and use a ternary vector of size $|\mathcal{F}|$ to determine which literals are observed at some specific state. The fact observations across all states in \mathcal{S} can be encoded as a ternary matrix of size $|\mathcal{F}| \times |\mathcal{S}|$, with the ternary value of entry (i, j) indicating whether $\text{holds}(F_i, S_j) \in \mathcal{C}$ holds, $\text{holds}(\neg F_i, S_j) \in \mathcal{C}$ holds, or neither. Such a matrix has an immediate representation as a non-negative integer (e.g., by reading its entries top to bottom and left to right, as a single ternary number). This integer can, then, be taken to encode the passive plot of the narrative.

Definition 8 (Canonical Narrative Index). Consider a narrative $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ w.r.t. some domain. The **canonical index** of $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ is a quadruple $\langle \lambda, \sigma, \alpha, \pi \rangle$ determined as above, where λ is the **temporal length** of $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$, σ is the **temporal structure** of $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$, α is the **active plot** of $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$, and π is the **passive plot** of $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$. The **canonical order** of a set of narratives \mathcal{N} is the order of narratives by increasing canonical index.

The canonical narrative index is a domain-independent index of narratives. Hence, the canonical indices of the narratives of any particular domain will not span all possible values. Some of the indices do not, in fact, correspond to narratives with respect to any domain. This happens exactly in those cases where σ represents a matrix encoding a relation \preceq^s with cycles. In the interest of naturalness, we have chosen not to exclude such cases from having an index, although excluding them would be possible by a slightly more involved definition of σ (where the matrix entries would not simply be read as a binary number, but where instead the matrix would be given a number according to its order among all matrices of the same size that encode an acyclic relation).

Since narratives can be indexed, it is now meaningful to ask to generate narratives of a domain in increasing order of their canonical index. We next show that this is possible.

Theorem 4 (Narrative Generation in Canonical Order). Consider a domain \mathcal{D} over a time-line $\langle \mathcal{T}, \preceq^t \rangle$. Then, there exists an algorithm that given \mathcal{D} , enumerates all narratives w.r.t. \mathcal{D} in their canonical order.

Proof. Note first that the claimed algorithm cannot terminate, since narratives of larger temporal length can always be constructed. Hence, we shall establish that the claimed algorithm eventually generates any given narrative w.r.t. \mathcal{D} , and that generated narratives appear in their canonical order.

The algorithm is the following: Consider a canonical index $\langle \lambda, \sigma, \alpha, \pi \rangle$, starting from $\langle 0, 0, 0, 0 \rangle$. Check to see if σ encodes a cyclic relation (e.g., by raising the corresponding matrix to all powers up to λ , and checking to see if the diagonal contains non-zero entries), and if not, then generate the corresponding discourse $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$. By Theorem 1, decide whether $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ is a narrative w.r.t. \mathcal{D} , and if so, output $\langle \lambda, \sigma, \alpha, \pi \rangle$ and $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$. In all cases, proceed to the next value of the canonical index, and repeat the process.

The only point that needs to be established is that indeed the algorithm can enumerate all indices in increasing order. But this is true by the following observation: Once λ is fixed, all remaining parameters of the index have a definite upper-bound. Consider, therefore, all possible values to the parameters σ, α, π in increasing order, before increasing λ by one, and repeating the process. \square

Theorem 4 implies, in particular, that the set of all possible narratives is computably enumerable. It is now straightforward to also compute the index of a given narrative.

Theorem 5 (Canonical Narrative Index Identification). Consider a domain \mathcal{D} over a time-line $\langle \mathcal{T}, \preceq^t \rangle$, and a narrative $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ w.r.t. \mathcal{D} . Then, there exists an algorithm that given $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ and \mathcal{D} , terminates and returns the canonical index of $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$.

Proof. The algorithm is the following: Run the algorithm that enumerates narratives w.r.t. \mathcal{D} in their canonical order, whose existence is guaranteed by Theorem 4. The algorithm will eventually output $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ along with its canonical index $\langle \lambda, \sigma, \alpha, \pi \rangle$. Return $\langle \lambda, \sigma, \alpha, \pi \rangle$, and terminate. \square

The definitions and results above can be extended to the case of default domains. From a technical point of view, this

extension is immediate, since a default domain is, simply, a collection of sub-domains with a preference order. The canonical index $\langle \lambda, \sigma, \alpha, \pi \rangle$ can be extended with an extra parameter, to account for the extra dimension that default domains introduce, namely preference. The complication that arises is a conceptual one, and shows itself once we attempt to determine the placement of this extra parameter in the index. If, on the one hand, it is placed before λ , and since λ can be infinitely large, then trying to generate narratives of a default domain in their canonical order will generate only preferred narratives and will never generate less preferred narratives. If, on the other hand, it is placed after λ , then certain less preferred narratives will be generated before certain other preferred narratives (although the latter narratives will be of larger temporal length). It is unclear to us which approach is conceptually more appropriate. We would be inclined to suggest the latter one, as it is guaranteed to eventually generate every narrative, preferred ones or less preferred ones, in increasing order of temporal length. If certain less preferred narratives need to be discarded, this can be easily done, since the index of each narrative will indicate also how preferred each narrative is.

Epilogue and Possible Sequels

A formal framework was presented, within which computational questions pertaining to narrative understanding can be studied. Within this framework certain fundamental computational properties of narrative have been established: the decidability of narrative recognition; the existence of preferred narrative; the ordering of narratives in preference order; the computability of a narrative canonical index; and the computable enumerability of the set of all possible narratives.

The two main premises of our framework, the necessity of a narrative context and the intimate relation of narrative and commonsense knowledge, are shared with previous approaches (see, e.g., [Verheij, 2009; Mueller, 2009]). Our framework especially echoes the view of Mueller [2009] on commonsense knowledge being the ultimate vehicle through which narrative can be understood and reasoned with. Our study, however, diverges from previous work in that it places emphasis on the formal implications that these premises have on the computability of various properties of interest for narrative. Our study also diverges from traditional theories of narrative as discussed by Altman [2008], in that it does not impose stringent requirements on what constitutes a narrative. Instead, it takes a sufficiently abstract view of narrative so as to not a priori exclude certain realizations or extensions that may capture more refined and intricate aspects of narrative, and so that the positive computability results be applicable to other more specific definitions of narrative.

Since narrative can be viewed as a lopsided conversation, a theory of narrative should be able to take a clear stance on the celebrated conversational maxims of Grice [1975]. According to these maxims, then, a narrative should contribute only what is: (i) believed to be true, (ii) sufficiently informative, (iii) relevant, and (iv) clear and unambiguous. Within our framework, (i) is enforced by the semantics, which insists that a narrative be true with respect to a context of

knowledge and beliefs, whereas (iv) is enforced by the syntax, which insists that a narrative be encoded in a logic-based unambiguous language. (ii) and (iii) are neither enforced nor precluded by our framework, since either case would require that the expectations of the narrative listeners be somehow encoded and taken into account. Although encoding expectations could presumably be done within a domain (or a natural generalization thereof), taking these expectations into account is not readily supported by the existing framework, suggesting, hence, an interesting direction for future work.

Several other extensions of this work are possible. Certain scholars suggest that having characters is an important ingredient of narrative (see, e.g., [Altman, 2008]). This aspect could be accommodated easily in our framework, and could build on the syntax of existing multi-agent frameworks (see, e.g., [Michael, Parkes, and Pfeffer, 2010]) to make explicit which agent is responsible for the execution of each action.

Rather intriguing would be an extension to accommodate for narratives in fictional worlds. When reading a fairy tale, an agent is expected to consider as context the knowledge that the fairy tale itself provides about the fictional world; in a precise sense, the narrative contains the constraints that it is expected to satisfy, and with respect to which it is to be understood. In case the fictional world knowledge conflicts with real world knowledge (which we assume is encoded in a domain available to the agent), the narrative-derived knowledge should take precedence. Our framework easily supports such a treatment without essential modifications.

An extension that could help in making our definition of narrative more culturally diverse would be that of accommodating narratives that do not satisfy the law of excluded middle. Such narratives appear in certain Eastern cultures, and correspond to singular discourses in our framework.

Among numerous other problems that one could attempt to formalize and solve through an extension of the presented framework, we mention only the problem of *narrative revision*, which we define to be the problem of (minimally) revising a given discourse so that it becomes a (preferred) narrative. Although this problem can be casted and studied by employing techniques similar to those that we have already used for defining default domains, space constraints have prevented us from discussing this problem further herein.

Beyond the framework extensions that we have discussed above, it would be beneficial to develop an actual system that implements the proposed framework. Such a system would serve as a vehicle to explore empirically — against, perhaps, a human gold standard — the abilities and limitations of the framework in recognizing and generating narratives, and help, in this manner, inform the extensions of the framework that would be most fruitful. For such a system to be effective, numerous issues need to be addressed. Mapping from natural language to logic can be done by employing techniques such as those found in [Bos and Markert, 2005] or [Michael and Valiant, 2008]. Computing models in default domains can be done by employing techniques such as those found in [Reiter, 1980; Michael and Kakas, 2009; Kakas, Michael, and Miller, in print]. Identifying preferred narratives can be done through the use of a typical semantics of argumentation [Bondarenko et al., 1997], and existing

systems that implement such a semantics can be employed.

The focus of this work has been in establishing the computability of certain central properties of narrative. Once this is done for a sufficiently rich set of properties, the natural next step is to study complexity. How can recognition, generation, indexing, and other processes be done *efficiently*. As narrative understanding is tied to computing domain models, existing complexity results for the latter problem could be informative, and existing efficient implementations of model computation could be used as subroutines of narrative understanding [Dimopoulos, Kakas, and Michael, 2004].

A research direction complementary to the ones above is that of identifying the domains within which narratives are to be interpreted. Although approaches to manually address this problem have been considered in the past [Lenat, 1995], we argue that a more viable, robust, and automated approach may come through the induction of domains: the process by means of which a given set of discourses that are assumed to be narratives with respect to a fixed, but unknown, target domain, is mapped to an approximation of the target domain.

From one point of view, a statement in a narrative given during the domain induction phase, can be thought of as offering a glimpse of some underlying hidden reality, about which the learner is expected to induce static commonsense knowledge. Learning frameworks that have been developed to deal with this problem from a purely logic-based point of view show that such an induction problem can be addressed to a certain extent. In the case where narratives are encoded in natural language text, the induction problem is directly relevant to the problem of Recognizing Textual Entailment [Dagan, Glickman, and Magnini, 2005], and existing work on that problem may be brought to bear [Michael, 2009].

From a second point of view, the sequence of statements within a narrative given during the domain induction phase, can be thought of as a set of snapshots of the evolution of the environment, about which the learner is expected to induce the underlying dynamics. This induction problem is known to be intractable even under certain simplifying assumptions [Michael, 2007]. It is our belief, however, that the problem merits further investigation, so that conditions are found under which the induction task is both useful and tractable.

It is worth emphasizing that, in the context of this work, success in the task of identifying domains is important only to the extent that *computerized* methods for dealing with narratives are to be developed. In particular, the appropriateness of our framework in offering a descriptive or prescriptive explanation of how *humans* deal with narratives, does not hinge upon the outcome of that task, since humans may be assumed to have identified through evolution or learning a domain that encodes their beliefs and knowledge. This explanatory aspect of our framework would be more appropriately tested by means of psychological or linguistic studies.

Opportunities for applying narrative understanding are certainly plentiful, and making progress along the discussed directions would bring us a step closer to realizing these opportunities. Concrete applications range from giving or understanding route instructions (thinking of the instructions as a narrative with respect to the context of an applicable map), to determining which among two legal arguments to

accept [Verheij, 2009] (thinking of the arguments as narratives with respect to the context of the law, and accepting the most preferred of the two). In fact, as Altman [2008] notes, “*Virtually any situation can be invested with [those] characteristics [necessary to] perform the narrational function*” and, hence, virtually anything can be thought of as a narrative. . . even the contents of a paper like this one!

Acknowledgments

The author would like to thank the anonymous Narrative’10 reviewers for some useful pointers.

References

- Altman, R. 2008. *A Theory of Narrative*. Columbia University Press.
- Bondarenko, A.; Dung, P.; Kowalski, R.; and Toni, F. 1997. An Abstract Argumentation-Theoretic Approach to Default Reasoning. *Artificial Intelligence* 93(1–2):63–101.
- Bos, J., and Markert, K. 2005. Recognizing Textual Entailment with Logical Inference. In *HLT/EMNLP’05*.
- Dagan, I.; Glickman, O.; and Magnini, B. 2005. The PASCAL Recognizing Textual Entailment Challenge. In *RTE’05*.
- Dimopoulos, Y.; Kakas, A. C.; and Michael, L. 2004. Reasoning about Actions and Change in Answer Set Programming. In *LPNMR’04*.
- Grice, P. 1975. Logic and Conversation. In Cole, P., and Morgan, J. L., eds., *Speech Acts*, volume 3 of *Syntax and Semantics*. Academic Press. 43–58.
- Kakas, A. C.; Michael, L.; and Miller, R. in print. Modular-E: An Elaboration Tolerant Approach to the Ramification and Qualification Problems. *Artificial Intelligence*.
- Lenat, D. B. 1995. CYC: A Large-Scale Investment in Knowledge Infrastructure. *CACM* 38(11):33–38.
- Michael, L., and Kakas, A. C. 2009. Knowledge Qualification through Argumentation. In *LPNMR’09*.
- Michael, L., and Valiant, L. G. 2008. A First Experimental Demonstration of Massive Knowledge Infusion. In *KR’08*.
- Michael, L.; Parkes, D. C.; and Pfeffer, A. 2010. Specifying and Monitoring Economic Environments using Rights and Obligations. *Autonomous Agents and Multi-Agent Systems* 20(2):158–197.
- Michael, L. 2007. On the Learnability of Causal Domains: Inferring Temporal Reality from Appearances. In *Commonsense’07*.
- Michael, L. 2009. Reading Between the Lines. In *IJCAI’09*.
- Mueller, E. T. 2009. Story Understanding through Model Finding. In *Narrative’09*.
- Reiter, R. 1980. A Logic for Default Reasoning. *Artificial Intelligence* 13(1–2):81–132.
- Verheij, B. 2009. Argumentation Schemes, Stories & Legal Evidence: A Computational Perspective. In *Narrative’09*.