A Standard Model of Mind:
AAAI Technical Report FS-17-05

Narrative as a Fundamental Information Unit in the Mind

Carlos León
Department of Software Engineering and Artificial Intelligence
Computer Science Faculty
Universidad Complutense de Madrid
28040 Madrid, Spain

Abstract

A standard model of the mind requires not only an architectural description of its components, but also a robust, general way to represent the information to be managed by it. With this purpose in mind, we explore to address knowledge representation in a standard model of the mind by studying how narrative, as an assumed fundamental way of structuring knowledge, can provide a valid framework to represent part of the information stored in the declarative memory. This paper studies how this information could be represented in terms of specific relations and how a standard cognitive architecture could make use of narrative information.

Introduction

The objective of a standard model of the mind is to provide a consensus for a computational description of human-like minds. Inspired by the standard model of physics, and based on existing cognitive architectures, Laird et al. recently proposed a seminal version of such a model (Laird, Lebiere, and Rosenbloom 2017).

For this standard model of the mind to be appealing to the community and useful in general, it should address a model of how each process works besides the pure architectural layout. In particular, a complete standard model should include a knowledge model, that is, a description of what the individual pieces of information are. While the specific representation could probably be implementation-dependent, the kind of information that the model manages and how that information is connected are aspects that a standard model of the mind could propose.

Research on knowledge representation has produced a diversity of models and formalisms along the years. In particular, semantic information has been represented with rules, ontologies, frames and others. All these formalisms have been relatively successful for some applications, but have failed to provide a commonly accepted model for representing knowledge. Most of the times, the bottleneck was not the formalism on which represent the information, but the knowledge model, that is, the facts, entities and relations that must me represented in the formal language at hand.

n this paper we seek to propose a number of basic aspects and relations between them for representing a subset of the declarative knowledge in a cognitive agent. We do not seek to provide a formalism to represent knowledge in a standard model of the mind, not to propose a method to solve knowledge representation in general. The objective is to explore one family of declarative knowledge (narrative knowledge) as a potential framework to enrich data in a cognitive system.

This declarative information is meant to be used in the declarative modules of a cognitive architecture. The idea is to root these relations in a coherent model that providing both flexible representation mechanisms and straightforward semantics for designing and understanding the underlying processes in a cognitive agent.

In particular, we propose to use narrative structure to build that framework. When seen as a model of representing information, narrative can provide a flexible and manageable model. In recent years, the understanding of narrative has evolved from the assumption that it is only a literary object and currently several authors support the idea that the ubiquity of narrative in all known human cultures is due to the very fundamental role it plays in human cognition.

This insight has led to the formulation of the so-called *narrative hypothesis*, stating that narrative "is not only a prominent form of human communication but also a fundamental way to represent knowledge" (Szilas 2015). Other authors agree on the role of narrative structures in cognition. (Bruner 1991). Schank and Abelson suggested that narrative can be used not only to report an episode, but also to structure knowledge (Schank and Abelson 1977). The so-called scripts, sort units of declarative steps describing a process, are proposed as a fundamental unit to represent information. Bruner argues that experience and memory happen "in the form of narrative". Herman proposes a *Story Logic* that assumes that narrative provides a framework to reason about experience (Herman 2002), and León proposes a list of fundamental relations for knowledge representation with narrative (León 2016). According to this, framing knowledge representation in terms of narrative features is not only possible but also useful.

This study assumes the *narrative hypothesis* and its relation with general models of cognition with the intention of supporting the inclusion of narrative as a main unit of information in the understanding of general cognition.

The current research as a very clear focus on artificial cognitive systems as computational processes, and it is strongly grounded on knowledge representation in artificial intelligence. While some of the assumptions and the implications might appeal to less practical fields intersecting cognitive science, most of the claims in this paper must be understood from this perspective.

It is therefore important to note that all the assumptions here, while inspired by biological cognitive systems, are meant to be understood as functional models of cognition, and as a potential part of a standard model. No claim is made about the plausibility of these process in biological agents.

Narrative as Structured Information

Although the analysis of narrative dates back to the very dawn of western philosophy with Artistotle (Baumrin 1986), a more systematic study of the underlying structure of narrative can be assumed to start with modern narratology in the XIX century. As a discipline, narratology provides several, not necessarily compatible definitions of what a narrative is. Some of the earliest ones are focused on the structural aspects of literary narrative (Propp 1928; Barthes 1977; Genette 1979). However, part of the narratology community has recently shifted towards cognitive aspects of narrative, thus providing definitions of it that are easier to include in cognitive architectures (Herman 2002; Abbott 2008).

The literature addresses many aspects of narrative, and the definitions vary depending on the author. Some of the aspects are to some extent shared by all definitions, like the role of events, causality and overall structure (Abbott 2008). Providing a widely accepted definition of narrative is out of scope. We will assume the next definition of narrative, that tries to capture the intuitive idea of narrative as a cognitive process and not only as a literary artefact:

> A narrative is a piece of information whose components are connected by *state*, *action*, *causality*, *time*, *location* and *agency*, and that shows overall coherence and completeness at a specific level of abstraction.

This description tries to capture two major components of narrative: 1) the existence of a specific set of features that link facts together as a narrative, and 2) the existence of a global coherence as a single object in a narrative.

In order to be more specific, we will be using the term *narrative unit* to make reference to each subnarrative in an overall narrative. For instance, the narrative describing how to create a computer program would be a narrative unit, and the subnarrative describing how to write a function would another narrative unit, used by the former.

The rest of the paper uses this as a working definition for structuring the study. While it is meant to be general enough to provide coverage to a wide range of definitions of narrative, some particular aspects are difficult to capture. This is discussed later on in the paper.

From the Physical World to Narrative Units

In this work we are assuming that narrative is a fundamental aspect of cognition, and not only a particular cognitive process. We believe this explains why narrative is so crucial for human communication and why, as a phenomenon, it is observable in all human societies.

As part of a physical world constrained by time and space, all perceptions of the environment happen along time and in specific places. We hypothesize that the abstraction of these physical constraints form the basic narrative relations: location and time as fundamental physical magnitudes, state and action as change of state, and agency and causality, are abstractions of the basic observable physical phenomena from the point of view of a cognitive agent. Following the idea that narrative basically captures physic phenomena and helps to describe them through symbolic relations, we believe it is a natural model for framing knowledge representation in artificial cognitive systems.

We take this idea to leverage the role of narrative, and propose *episodic perception*. Through episodic perception, information is connected and structured as a narrative as part of the process of perception, and not only as a post-hoc process. That is, the acquisition process produces symbols structured as narratives, in parallel with raw symbol perception. The fundamental relations hold for all declarative information that can be structured in narrative terms.

The model of episodic perception does not try to replace accepted models of acquisition, it only provides an explanation of why narrative happens in such a natural way in humans. Next sections are devoted to the explanation of this functional model.

Narrative Units from Episodic Perception

As detailed in the previous section, artificial cognitive systems can be modelled as systems in which perception is episodic. In order to structure information for storage or learning and retrieval in declarative memory, the atomic components of the episode must be linked together in a single element. The semantics of these links must be stable across all modules for a cognitive architecture to be able to use it.

The first step to perform a semantic processing in a cognitive architecture is to identify raw perception and aggregate the stimuli into entities. How this is done is outside the scope of this research, but it is assumed that this transformation from the external world into symbols must occur prior to any semantic processing.

The second step is to structure symbolic information into narrative units. According to Chatman's work, narratives are formed by the *kernels* and the *satellites* (Chatman 1978). Kernels are the most salient, prominent parts of a narrative, while satellites are pieces of knowledge connected to the kernel that complement it by providing context, causality or other additional information. Following this, we propose that narrative units in a general model of cognition must be structured in terms of its kernels and satellites, that is, to identify these most salient aspects and build narrative units using the kernels as roots. A narrative unit would therefore have one single main *kernel*, and would have its satellites connected by other narrative relations, as depicted if fig. 1.

Narrative facts are the core elements conveying the basic meaning. The narrative kernels must therefore be narrative facts (although not all narrative facts are narrative kernels). These facts hold information about *states* and *actions*:

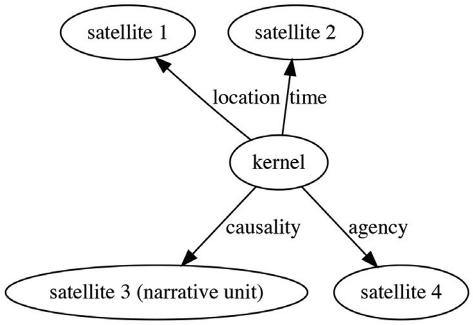

Figure 1: Graphical representation of kernels and satellites in a narrative unit.

State States are pieces of information that represent facts that hold for the current time slice. For instance, "John is awake" or things that always hold like "Isaac Newton was a scientist".

Actions Actions represent changes of state: "waking up", or "grabbing a glass". Actions represent narrative evolution, and are key in chaining sequential states together.

The differentiation between these two types in the model is intended to provide coverage to a number of reasoning processes in the way that many artificial intelligence algorithms do. The separation is therefore motivated by practical reasons. An alternative model could merge states and actions since actions and states can be both represented in terms of each other. For instance, "move" could be modelled as an action (from position x to position $x+1$) or a state (the current property of agent a is "moving"). This alternative, however, is not considered strictly necessary for the model, and taking states and actions together, like a single type of narrative facts, is probably useful.

In the proposed model, both states and actions are connected to the rest of the information by the fundamental narrative relations that this model addresses. A list of relations for narrative knowledge representation in cognitive architectures has been proposed before (León 2016). That list is intended to provide a formalism for representing narratives, but has anyway inspired the relations proposed in this model. These relations intend to address fundamental narrative aspects and not necessarily formal knowledge representation:

Time Time represents physical time, either in an absolute (*4 seconds*) or relative way (*before*). Time is fundamental in narrative, and models many physical constraints.

Location Location models the physical space. Like time, this could be absolute (*in the park*) or relative (*behind the building*).

Causality Causality is key in reasoning, and links the effect of applying actions on subsequent states. While time and location link narrative kernels to objects, causality links narrative units.

Agency Agency is fundamental in how narrative units are grouped. Narratives are constructed by focusing on the evolution of agents.

Narrative units, as high level structures over symbolic information in a cognitive agent, can be modelled according to these relations. Specific formal semantics are assumed to be dependent of the implementation, but it is hypothesized that the proposed relations are among the most fundamental ones.

An implementation of the standard model would still have to provide a knowledge representation formalism for representing actions, states, agents, time and so on.

Storing Narrative Units

It is common for prominent cognitive architectures to divide memory into procedural and declarative (Laird 2012; Anderson 2007; Rosenbloom, Demski, and Ustun 2016). Additionally, declarative memory is usually divided into semantic and episodic memories, either structurally or at least functionally in some cases.

In this section we will argue how storage (learning) and retrieval of declarative information can benefit from the framing of information as narrative units. We will try to show how the narrative properties themselves potentially represent strong associations between elements in the working, episodic and semantic memory.

Narrative Units in Working Memory

Working memory (Baddeley 2012) plays a fundamental role in many cognitive architectures. These models usually place working memory as a central module interacting between input/output processes and as a container for buffers managing information from other modules (memory). In particular, the standard model proposed by Laird et al. (Laird, Lebiere, and Rosenbloom 2017) assumes that communication between processes mostly occur in buffers in the working memory.

Following this, we propose an active process happening during perception in which symbolic input is structured as a narrative unit. That is, each perception sample is received in the working memory and, in parallel, buffered in this *episodic buffer* until a narrative unit is created, and then provided to the working memory as a piece of information. This schema is depicted in fig. 2.

External perceived information feeds both the episodic buffer and other possible buffers of the working memory in parallel. The episodic buffer then provides both the working memory and the long-term memory with narrative units, whose underlying relations can enrich the storage or learning process in the long-term memory by linking the perceived narrative units with other stored data by agency, causality and so on.

In our model, the active process of the episodic buffer can use content from the declarative long-term memory in order to build the narrative units. Both the semantic memory and the episodic memory potentially store content needed when aggregating symbolic information into narratives.

Storing Narrative Units in Episodic Memory

Episodic memory is the natural storage of narrative units interconnected by the previously described relations. By promoting narrative structure into a general way of framing

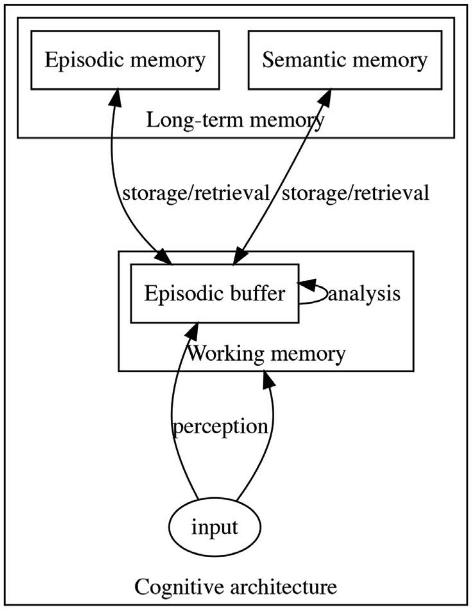

Figure 2: Episodic buffer and its relation with the working and long term memories in a cognitive architecture.

knowledge, episodic memory becomes a highly connected graph.

These connections happen, at least, at three different levels:

- Relations between narrative units and its satellites (specific agents, time, locations, etcetera). This symbolic relations would be directly or indirectly linked to its corresponding semantics in the semantic memory.
- Sequential information of states and actions, forming episodes. These episodes would be enriched by causality and time. Additionally, the narrative information would permit identifying proper high level narrative units in these episodes, potentially chunking them into individual, well delimited narratives.
- Hierarchical relations of episodes. Episodes could be composed by identifiable sub-episodes, and script-like information in the semantic rules (instructions about processes) could be related.

As it will be explained later on, this semantic structure creates a graph whose now explicit information can be used by a retrieval process.

Storing Narrative Units in Semantic Memory

Narratives, according to the working definition in use, can provide a common vocabulary to inform the grounding of semantics in the semantic memory. According to our definition, narrative structures in cognition abstract and connect information perceive from the physical environment.

Schank et al. proposed that declarative processes can be represented in the form of scripts (Schank and Abelson 1977). In this model, scripts, as abstract steps than describe

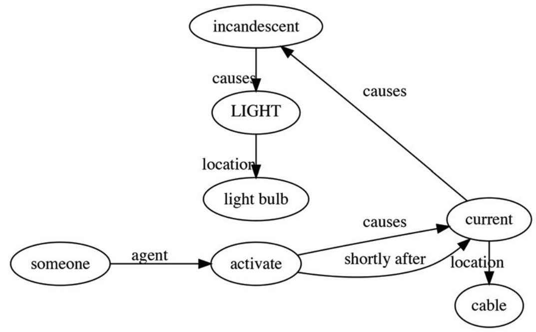

Figure 3: Example of how to encode semantic knowledge as a narrative.

common situations, are assumed to be part of what is stored in the semantic memory.

By abstracting a number of instances, a script can be created from episodic information. Given that we assume episodic information to be enriched with narrative relations, the corresponding script could also be represented, at least partially, in these terms, sharing the same vocabulary. The relations would still represent the same semantics and they could be used by the cognitive system to identify similar properties or structures between episodic and semantic information, therefore connecting the semantics between a script (information about how to do something) and one or several instances of that rule.

Hypothetically, grounding information in narrative can also make it possible to represent general rules like "electricity is used for light bulbs" by encoding it as a narrative. The fact will be then assumed to be a specific rendering of the short script describing how someone activates (agency) the switch the electric current. This causes a current (causality) after a short time (time) in the cable (location). The light bulb them becomes incandescent, which causes the light. The light, in this case, would be the nucleus of this narrative unit.

This narrative unit is graphically depicted in figure fig. 3.

Retrieving Narrative Units

By applying our model of narrative information in the working memory, the episodic memory and the semantic memory, the three modules can share a common vocabulary with specific semantics. Additionally, since the semantics represent basic concepts assumed to be abstractions of physical phenomena, the amount of common information shared between units is assumed to be relatively high. This would hypothetically produce a high level of interrelatedness, which can be naturally used to model association.

Based on this potentially strong associativity characteristics, retrieving past episodes can trigger the activation of other episodes by means of narrative properties like location or time. Scripts, framed with the same relations, could be activated to provide the semantics that contextualize an episode.

Retrieval of script-like information stored in the semantic memory can also make use of the general narrative relations that the proposed model provides. For instance, and using the famous example, having dinner in a restaurant can be represented not only by one specific script, but by many other hierarchically connected: information about asking for table, a script of how to order food, how to pay, etc. Location, time, causality and agency are general relations that apply in a vast diversity of rules. Analogously, the model of narrative units would seamlessly activate relevant instances of the script by identifying related places, times or agents.

Moreover, the instances and the rules could differ to some extent. The difference between the narrative unit that has been experienced and the one that was expected because of the information in the script can trigger a cycle of revision in the cognitive architecture.

Other Forms of Knowledge

As an example, specific facts like "Paris is the capital of France" does not necessarily need narrative relations to be represented in a cognitive system. However, the grounding of what a capital *means* can be carried out in terms of narrative: the kind of things that happen in a capital, the relation between a capital and a country, and so on.

As such, narrative units are not a replacement of widely used knowledge structures in artificial cognitive systems, but a natural complement to them. We argue that most basic, factual information can be grounded in narrative units, since narrative units are basically abstractions over physical phenomena. If one wanted to take the model to the limit, even abstract terms like mathematical abstractions could be defined in terms of the processes that humans use to apply them. This, however, would probably be unnecessary, and such an application of the model is not the purpose of this proposal.

While the presented theory strictly applies to a specific part of the memory in the standard model of the mind (declarative), the underlying hypothesis could also have implication for other forms of knowledge. For instance, non-declarative information like face images and its corresponding pattern matching for identifying a known face is usually modelled with statistical methods with relatively high level of success. The narrative hypothesis and the current framework could potentially be applied to these by finding out *narrative relations* between these non-declarative objects.

Conjecturally, one could think of a statistical learning process by which not only the points in an array image are fed as inputs, but also their potential narratively-related units like other images, sounds, even active declarative units present in the working memory. This would help to model associative memory by grounding associations precisely in these physical-narrative relations. Conversely, non-declarative knowledge could *provide* narrative relations. The output of a speech recognition process could yield not only the words as text or symbols, but also the agency (knowing "who" said something).

This potential use of narrative units in non-declarative knowledge roots in the influence of physics in the model. Physics, either described as symbols and equations or as non-symbolic relations in connectionist networks, is assumed to be the common ground on which cognition emerges. The presented theory of narrative as an information unit tries to frame this for the standard model.

Discussion

The current model does not address a number of aspects that are not purely narrative but play a role in knowledge representation. For instance, abstraction and quantification in semantic rules (to represent things like "all the inhabitants of Asia") is not addressed, although the model is compatible with them. In this sense, the model of narrative units is meant to provide a standard model of the mind with a high level, stable framework to enrich knowledge representation with relations that are assumed to hold for most pieces of information. The current model is not meant to substitute other models or even to enforce a way of representing knowledge in cognitive systems.

The model has been presented in terms of active processes (episodic perception, episodic buffer), basic pieces of narrative information (narrative facts, kernel and satellites) and relations (location, time, agency and causality), but no specific formalism for representing them in an artificial cognitive system has been proposed. This has been done intentionally, since current cognitive architectures, as the best instances of what can become reference implementations of the standard model, implement their own formalisms. However, it is conjectured that an implementation of these relations is, if not already implicitly present in all of them, directly implementable.

Finally, it is important to understand this study in the context of our working definition of a narrative. As previously introduced, there are a diversity definitions of narrative, and while the one that we propose is intended to provide coverage to a number of them, there is no doubt that it will not perfectly fit all others. As such, we believe the description is more important that the term.

In this regard, while the present proposal addresses narrative as a way to structure knowledge, it is important to keep in mind that we are assuming a meaning of narrative that do not necessarily have to match the common use of the word. In order to prevent conflict with other valid definitions of narrative, we strictly focus on the one concrete definition previously provided. Our understanding of narrative is tightly related with the eventful, related structure of symbolic knowledge in specific terms, and has little to do with any literary or even communicative aspect. The term is just considered to be appropriate to convey the meaning and to relate it with existing literature.

Conclusions and Future Work

A standard model of the mind must address structural and functional aspects of the mind. According to the *narrative hypothesis*, narrative structures can provide a robust framework for representing knowledge in cognitive system. Following this hypothesis, narrative should be addressed in a standard model of the mind. In order to provide a theoretical way to include narrative, this paper has suggested a model of how to use narrative structures in general cognitive architectures, both in terms of the relations that connect information as narrative units and the potential structural role of the modules handling these relations in cognitive architectures.

The model has been presented with the purpose of serving as an initial idea to identify common aspects of knowledge that must be addressed in the standard model as narrative. Even if the model results useful to express common processes, it is well possible that not all of it is valid, or that there are still some fundamental pieces missing. In any case, the model by itself is not of practical use. Instead, it must be tested in real cognitive architectures to test both its validity and the potential benefits that have been hypothesized.

One aspect that has not been studied in detail is the interpretation of raw stimuli as symbolic relations. This is usually a challenge for most aspects of knowledge representation, and the way it affects the creation of the declarative relations of narrative units is not essentially different. Aspects close to measurable aspects in the physical world, like location or time, could have reasonable approximations, but identifying causality or agency would require much more complex processes.

Acknowledgments

This work has been supported by the IDiLyCo project (TIN2015-66655-R) funded by the Spanish Ministry of Economy, Industry and Competitiveness; and by the projects WHIM 611560 and PROSECCO 600653 funded by the European Commission, Framework Program 7, the ICT theme, and the Future and Emerging Technologies FET program.

References

Abbott, H. P. 2008. The Cambridge Introduction to Narrative. *Assessing Writing* 252.

Anderson, J. R. 2007. *How Can the Human Mind Occur in the Physical Universe?*

Baddeley, A. 2012. Working Memory: Theories, Models, and Controversies. *Psychology* 63(1):1–29.

Barthes, R. 1977. The Death of the Author. In *Image, Music, Text*. New York, NY, USA: Hill and Wang.

Baumrin, B. H. 1986. The complete works of Aristotle: The revised Oxford translation.

Bruner, J. 1991. The Narrative Construction of Reality. *Critical Inquiry* 18(1):1.

Chatman, S. 1978. *Story and Discourse*. Cornell University Press.

Genette, G. 1979. *Narrative Discourse: An Essay in Method*. Cornell University Press.

Herman, D. 2002. *Story Logic: Problems and Possibilities of Narrative*.

Laird, J. E.; Lebiere, C.; and Rosenbloom, P. S. 2017. A Standard Model for the Mind: Toward a Common Computational Framework across Artificial Intelligence, Cognitive Science, Neuroscience, and Robotics. *AI Magazine (special issue on The Cognitive System Paradigm: A New Thrust to Attain Human-Level AI).*

Laird, J. E. 2012. *The Soar Cognitive Architecture*, volume 171.

León, C. 2016. An architecture of narrative memory. *Biologically Inspired Cognitive Architectures* 16:19–33.

Propp, V. 1928. *Morphology of the Folk Tale*. Akademija, Leningrad.

Rosenbloom, P. S.; Demski, A.; and Ustun, V. 2016. The Sigma Cognitive Architecture and System: Towards Functionally Elegant Grand Unification. *Journal of Artificial General Intelligence* 0(0):1–103.

Schank, R., and Abelson, R. 1977. *Scripts, Plans, Goals and Understanding: an Inquiry into Human Knowledge Structures*. Hillsdale, NJ: L. Erlbaum.

Szilas, N. 2015. Towards Narrative-Based Knowledge Representation in Cognitive Systems. In *Proceedings of the 6th Workshop on Computational Models of Narrative (CMN'15)*, 133–141.