Representing and Managing Narratives in a Computer-Suitable Form

Gian Piero Zarri

University Paris-Est, Créteil, Val de Marne (UPEC), LiSSi Laboratory
120-122 rue Paul Armangot, 94400 Vitry-sur-Seine, France
zarri@noos.fr, gian-piero.zarri@u-pec.fr

Abstract

Narratives can be defined informally as a "spatio-temporally bounded stream of elementary events". To make this sort of definition more computationally useful we introduce, firstly, some pragmatic criteria for recognizing highly ambiguous entities like the "elementary events" and for linking these events together into complete narratives. We raise then the problem of how to concretely represent elementary events and narratives in computer-suitable form. We introduce then the main characteristics of a language, NKRL (Narrative Knowledge Representation Language), expressly specified and implemented for dealing with (non-fictional) narratives and temporal information. We conclude by showing briefly how this language can be used for questioning and for particularly complex inference operations.

Introduction

Narrative information concerns the account of some real-life or fictional story (a "narrative") involving concrete or imaginary characters; a narrative can be defined informally, see below, as a "stream of elementary events that is spatio-temporally limited". In this paper, we will try first to give some concreteness to this very fuzzy kind of definition, introducing, e.g., pragmatic but operationally sound criteria for defining ambiguous entities like the above "elementary events". Passing then to the problem of how representing events and narratives in a computer-usable form, we will introduce NKRL (Narrative Knowledge Representation Language), see Zarri (2005, 2009a, 2009b), a representation language and (fully implemented) computer science environment specifically created for dealing with challenging forms of narratives.

We can note that, in an NKRL context, we are mainly concerned with *non-fictional narratives*, like those typically embodied into corporate memory documents (memos, policy statements, reports, minutes etc.), news stories, normative and legal texts, medical records, many intelligence messages, surveillance videos, actuality photos for newspapers and magazines, material (text, image, video, sound…) for eLearning, Cultural Heritage material, etc. This choice corresponds only to practical constraints – to benefit, e.g., from the financial support of the European Commission (EC) – and nothing (apart from considerations of time, appropriateness, amount of code etc.) might prevent us from dealing in an NKRL mode with the whole "Gone with the wind" *fictional-narrative* novel.

In the following, we will deal first with the theoretical foundations, providing then the general framework of the NKRL approach. We will then supply some concrete details about this language/environment: this Section will include two sub-sections, dealing the first with the knowledge representation aspects and the second with the querying/inference operations. A short conclusion will close the paper.

Theoretical Foundations

Our understanding is that (fictional or non-fictional) narratives correspond essentially to the basic layer, the "*fabula* (a Latin word: fable, story, tale, play) *layer*", introduced by Mieke Bal (1997) in her seminal work on the structures of narrative phenomena. Accordingly, an (NKRL-like) narrative can be seen informally as a *series of logically and chronologically related events (a 'stream of elementary events') that describe the activities or the experiences of given characters*. From this basic definition and other work in a *narratology* context we can already infer some important characteristics of (fictional/non fictional) narratives, see also Zarri (2009b: 2-13):

- One of the features defining the *connected* character of the elementary events that make up the stream concerns the fact that these are *chronologically related*, i.e., narratives/complex events *extend over time*. This diachronic aspect of narratives (a narrative normally has a *beginning*, an *end* and some *form of development*) represents one of their most important characteristics.
- *Space* is also very important in the narrative domain, given that *the elementary events of the stream occur in well defined 'locations'*, real or imaginary ones. The connected elementary events that make up a narrative are then both *temporally and spatially bounded*.

- A simple chronological successions of elementary events that take place in given locations cannot, however, be defined as a narrative without some sort of *semantic coherence* and *uniqueness of the theme* that characterize the different elementary events of the stream. If this logical coherence is lacking, these events pertain to different narratives: a narrative can also be represented by a single elementary event.
- When the constitutive elementary events of a narrative are verbalized in NL terms, their coherence is normally expressed through syntactic/semantic constructions denoting causality, goal, indirect speech, co-ordination and subordination, etc. In this paper, we will systematically make use of the term *connectivity phenomena* to denote this sort of clues, i.e., to denote what, in a stream of elementary events, i) leads to a global meaning that goes beyond the simple addition of the meanings conveyed by the single events; ii) defines the influence of the context where a given event is used on the meaning of this individual event, or part of it.
- Eventually, narratives concern the behavior or the condition of some actors (persons, characters, personages, figures etc.). They try to attain a specific result, experience particular situations, manipulate some (concrete or abstract) materials, send or receive messages, buy, sell, deliver, etc. In short, *they have a specific role in the event* (in the stream of elementary events representing the global narrative). Note that these actors or characters are not necessarily human beings; we can have narratives concerning, e.g., the vicissitudes in the journey of a nuclear submarine (the actor, character etc.) or the various avatars in the life of a commercial product.

Defining, however, a narrative as a *"spatio-temporally bound stream of elementary events"* would be of a scarce concrete utility without being able to specify what an *elementary event* is. In an NKRL context, this point is also particularly important from a practical point of view given that, as we see later, *each elementary event is separately encoded making use of the NKRL knowledge representation tools*.

According then to a well-known Jaegwon Kim's definition, see Kim (1996), a "*monadic*" event – which can be considered as equivalent to an elementary event – is identified by a triple $[x, P, t]$ where x is an object that exemplifies the *n*-ary property or relation P at time t (where t can also be an interval of time). Monadic means that the *n*-ary property P is exemplified by a single object x at a time. To make reference to one of the recurrent examples in the theoretical discussions about events, "Brutus stabs Caesar", the Kimian interpretation of this event corresponds then to the representation of an individual x, Brutus who, at time t, is characterized by the property P exemplified by his stabbing of Caesar. Without entering now in the theoretical controversies raised by this sort of definition, see Zarri (2009b: 8-13) for some information in this context, we can note that this description is both:

- *interesting* because of the stress on the properties that characterize the main actor of the elementary (monadic) event, and the importance of time;
- *insufficient* with respect to the opportunity to make use practically of this kind of definition because, mainly, of the lack of details about the possible relationships between x and P.

A more complete, structured and 'useful' definition of elementary events has been introduced by Donald Davidson, see Davidson (1967); this is particularly popular in the linguistic domain. Davidson's definition centers the representation of an elementary event around the "*action verb*" characterizing the *global conceptual category* of the event more than – as in the Kimian approach – on the "*generalized properties*" of this event. In this way, the Davidsonian representation of "Brutus stabs Caesar" becomes: $\exists e.\mathrm{stab}(e, b, c)$, where e is an *event variable*. The global meaning of this formalism corresponds then to: "There is an event e such that e was a stabbing of Caesar (c) by Brutus (b)". Moreover, as emphasized above when we have listed some important characteristics of narratives, *roles have a particular importance in a narrative*. Our preferred formalism for the representation of *elementary* events is then the so-called *neo-Davidsonian* approach. The neo-Davidsonians see, e.g., Higginbotham (2000), Parson (1990), assume in fact that the event argument e above must be *the only argument of the (verbal) predicate*: this implies then, necessarily, *the introduction of thematic (functional) roles for expressing the relations between events and their participants*. The formalization of "Brutus stabs Caesar" becomes then now:

$$\exists e[\mathrm{stab}(e)\ \&\ \mathrm{agent}(e) = b\ \&\ \mathrm{object}(e) = c]\ . \qquad (1)$$

Apart from the theoretical implications, what expounded above is particularly important from a practical point of view because it supplies us with a *pragmatically useful and operational criterion for recognizing and isolating – in some way, for 'defining' – elementary events*. The criterion consists then in the identification, *within the description in natural language (NL) terms of the global stream representing a narrative/complex event, of a specific generalized natural language predicate*: this will then represent the core of a new elementary event. The predicate will correspond usually to a *verb* – to stick to the previous example, recognizing "to stab" as a verb in the NL chain "Brutus stabs Caesar" should be sufficient for signaling the presence of an elementary event – but, according to the neo-Davidsonian approach, this predicate can also, in case, correspond to a noun (…Jane's *amble* along the park…) or an adjective ("… *worth* several dollars…") *when these last have a predicative function*.

Of course, a drawback of this criterion concerns the fact that its utility is limited to the recognition of the elementary components of narratives/complex events that can be expressed in NL terms, while narratives are multimedia in essence – a photo representing President Obama addressing the Congress, or a short video showing three nice girls on a beach, are surely narrative documents (the first including only an elementary event) but they are not, of course, NL documents. A classical way of getting around this problem however exists, and it consists in annotating multimedia narratives in natural language.

Formal Representation of Narratives

Eq. 1 above – an *n*-ary form of representation – shows clearly that the now so popular W3C proposals like RDF(S), OWL or OWL 2 are, *at least in their standard format*, unable to supply a basis for representing elementary events on a computer. All the W3C representations are, in fact, based on the usual *attribute – value* model, where a property/attribute can only be a *binary relationship* linking two individuals or an individual and a value. The inadequateness of this approach to take into account *complex representational problems* like those linked with narratives, spatio-temporal information, events and complex events etc. is today widely recognized see, *inter alia*, (Hoekstra et al. 2006, Mizoguchi et al. 2007, Salguero, Delgado and Araque 2009, Liu et al. 2010).

Note that the argument often raised in a W3C context and stating that any representation making use of *n*-ary relations can be always converted to one making only use of binary relations *without any loss of expressiveness* is incorrect with respect to the last part of the sentence. It is true in fact that, from a pure formal point of view, any *n*-ary relationship with $n > 2$ can always be reduced, in a very simple way, to a set of binary relationships. However, this fact does not change at all the *intrinsic, semantic* n*-ary nature* of a simple statement like "Bill gives a book to Mary" that, to be understood, requires to be taken *in its entirety*. This means – see also Eq. 1 above – to make use of a semantic predicate of the GIVE type that introduces its *three* arguments, "Bill", "Mary" and "book" through *three* functional relationships (roles) like SUBJECT (or AGENT), BENEFICIARY and OBJECT, the whole *n*-ary construction being – this is the central point – *reified and necessarily managed as a coherent block at the same time*. Only in this way it will be possible to infer that, e.g., the above elementary event is linked, in the framework of a wider narrative, to another elementary event relating Mary's birthday; for more details see, e.g., Zarri (2005b: 14-21).

Several *true* *n*-ary models able to represent in a computer-usable way elementary events have been described in the literature, see Zarri (2009b: 22-33) for a review. The *n*-ary model used in NKRL to represent these events can be denoted as:

$$(\mathsf{L}_i\,(\mathsf{P}_j\,(\mathsf{R}_1\,\mathsf{a}_1)\,(\mathsf{R}_2\,\mathsf{a}_2)\,\ldots\,(\mathsf{R}_n\,\mathsf{a}_n)))\;, \qquad (2)$$

where L_i is the symbolic label identifying ("reifying") the particular *n*-ary structure (e.g., the global structure corresponding to the representation of the previous "John gives a book to Mary" example), P_j is a conceptual predicate, R_k is a generic functional role and a_k the corresponding predicate argument (e.g., the individuals JOHN_, MARY_ etc.). Note that each of the $(\mathsf{R}_i\ \mathsf{a}_i)$ cells of Eq. 2, *taken individually*, represents a *binary relationship* in the W3C (OWL, RDF…) languages style. The main point is here that, as already stated, *the whole conceptual structure represented by (2) must be considered globally*.

Similarities between neo-Davidsonian expressions for elementary events like that of Eq. 1 and the formal structure of Eq. 2 are evident. However, some important differences exist. To avoid both the typical ambiguities of natural language and possible combinatorial explosion problems – see the discussion in Zarri (2009b: 56-61) – both the (unique) conceptual predicate of Eq. 2 and the associated functional roles are *primitives*. Predicates P_j pertain in fact to the set {BEHAVE, EXIST, EXPERIENCE, MOVE, OWN, PRODUCE, RECEIVE}, and the functional roles R_k to the set {SUBJ(ect), OBJ(ect), SOURCE, BEN(e)F(iciary), MODAL(ity), TOPIC, CONTEXT}. Two special operators, date-1 and date-2 – that can be assimilated to functional roles – are used to introduce the temporal information associated with an elementary event: see, e.g., Zarri (2009b: 76-86, 194-201) for a detailed description of the *formal system* used in NKRL for the representation and management of temporal information. The NKRL representation of *specific elementary events* – that corresponds then to the *concrete instantiations* (called "*predicative occurrences*" in the NKRL's jargon) of general structures in the style of Eq. 2 – is then a sort of *canonical representation*.

Several predicative occurrences – denoted by their symbolic labels L_i and *representing formally a (possibly structured) set of elementary events* – can be associated within the scope of second order structures called "*binding occurrences*". These are, in practice, *labeled lists* formed of a *binding operator Bn with its arguments*. The *Bn* operators are those used in NKRL to represent the *connectivity phenomena*, see the previous Section, which guarantee the *global coherence of narrative/complex events*. They are: ALTERN(ative), COORD(ination), ENUM(eration), CAUSE, REFER(ence) – the "*weak causality*" operator, introducing two arguments where *the second is necessary but not sufficient to explain the first* – GOAL, MOTIV(ation) – the "*weak intentionality*" operator, where the *first argument is not necessary to carry out the second, which is however sufficient to explain the first* – COND(ition), see Zarri (2009b: 91-98). The general expression of a binding occurrence is then:

$$(Bn_k\ \mathsf{arg}_1\ \mathsf{arg}_2\ \ldots\ \mathsf{arg}_n)\;, \qquad (3)$$

Eq. 3 is particularly important in an NKRL context because it supplies also a *formal expression of the notion of narrative in agreement with the intuitive definition supplied at the beginning of this paper*. The arguments arg_i of Eq. 3 can, in fact, i) correspond *directly* to L_i labels – i.e., they can denote simply the presence of particular elementary events represented formally as predicative occurrences – or ii) *correspond recursively to new labeled lists in Eq. 3 format*. In the first case, the global narrative represents merely *a chronological stream of elementary events, temporally characterized, where all these events have the same logical/semantic weight* and the operator *Bn* corresponds to COORD (or ENUM/ALTERN). In the second case, we can suppose, e.g., that a given sequence of events – an Eq. 3 list of the COORD... type – represents the CAUSE of another sequence of events. The global representation of this narrative/complex event will then correspond to an Eq. 3 list labeled as CAUSE, having as arguments arg_1 the COORD... list including the elementary events at the origin of the narrative, and as arg_2 the COORD... list including the elementary events that represent together the consequence, see also the example at the end of the first sub-section below. What expounded above is in agreement with the remarks expressed by several authors – see Mani and Pustejovsky (2004) for example – about the possibility of *visualizing under tree form the global, formal expression of a narrative/complex event made up of several elementary events*.

A Short Description of NKRL

After having introduced, in the previous Sections, the general theoretical framework underpinning the NKRL approach to the narrative problems, we will now illustrate briefly some points concerning its concrete implementation – see Zarri (2009b) for a complete description.

An important point to be immediately noticed concerns the fact that NKRL has been developed from the beginning as a tool for expressing in the best way, at the deep level, the '*meaning*' of narratives, *without caring of all about the form these narratives can assume, at the surface level, in a particular natural language*. The problem of establishing some form of correspondence – e.g., in the Jackendoff's (1990) sense – between deep and surface levels obviously exists, *for practical reasons*, also in an NKRL context, but this problem is tackled *a posteriori and in a fully pragmatic way*, see Zarri (2009b: 246-248) for some information in this context.

This sort of *deep meaning approach* explains also why the similarities between NKRL and some 'linguistic' and 'surface' oriented tools like FrameNet (Johnson et al. 2001) and WordNet (Kipper et al. 2005) are *quite superficial*, even if these last tools can be obviously be of some utility from the practical point of view evoked before. The goals of NKRL and of the previously mentioned lexical resources are, eventually, totally different – e.g., the main, final aim of NKRL concerns chiefly the possibility of implementing *powerful inference operations on the contents of a formal knowledge base*, see Zarri (2005, 2009a) in this context. Also some formal similarities concerning the use of conceptual tools like "predicates" and "roles" in both NKRL and FrameNet/WordNet are more apparent then real. As already seen, predicates and roles in NKRL correspond to *deep primitives* and are very limited in number. Roles in FrameNet are highly idiosyncratic; Wordnet has 24 semantic roles and 94 semantic predicates, etc.

The Knowledge Representation Aspects

NKRL innovates with respect to the current ontological paradigms by adding to the usual ontologies of concepts an "ontology of (elementary) events", i.e., *a new sort of hierarchical organization where the nodes correspond to* n*-ary structures in the style of Eq. 2 above*. In the NKRL's jargon, these *n*-ary structures are called "templates" and the corresponding hierarchy – i.e., the ontology of elementary events – is called HTemp (hierarchy of templates). Templates can be conceived as the canonical, formal representation of generic classes of elementary events like "move a physical object", "be present in a place", "produce a service", "send/receive a message", etc.

Note that, in the NKRL environment, an *ontology of concepts* (according to the traditional meaning of these terms) not only exists, but it represents an essential component of this environment. This standard ontology is called HClass (*hierarchy of classes*): structurally and functionally, HClass is *not fundamentally different* from one of the ontologies that can be built up by using tools in a traditional (frame-like) Protégé style. An (extremely reduced) image of HClass is given in figure 1 – HClass includes now (September 2010) more than 7,000 concepts.

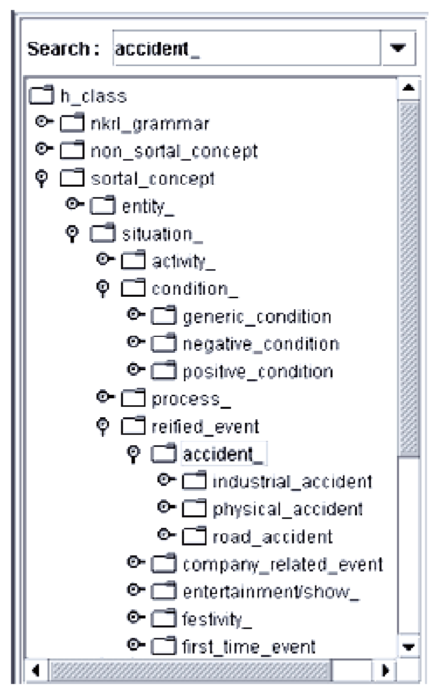

figure 1. Partial representation of HClass, the 'traditional' ontology of concepts.

When a *specific elementary event* pertaining to one of the general classes of elementary events denoted by templates must be represented, the corresponding template is *instantiated*, giving rise to a "*predicative occurrence*". To represent then a simple elementary event (*corresponding to the identification of the surface predicate "offer"*) like: "British Telecom will offer its customers a pay-as-you-go (payg) Internet service in autumn 1998", we must select firstly in the HTemp hierarchy the template corresponding to "supply a service to someone", represented in the upper part of table 1. This template is a specialization of the particular MOVE template corresponding to "transfer of resources to someone" – figure 2 below reproduces a fragment of the external organization of HTemp. In a template, the arguments of the predicate (the a_k terms in Eq. 2) are concretely represented *by variables with associated constraints*: these are expressed as HClass concepts or combinations of concepts, i.e., the two ontologies, HTemp and HClass, *are strictly intermingled*.

table 1. Deriving a predicative occurrence from a template.

```
name:   Move:TransferOfServiceToSomeone
father:   Move:TransferToSomeone
position:   4.11
natural language description: "Transfer or Supply a Service to
Someone"

MOVE      SUBJ             var1: [var2]
          OBJ              var3
          [SOURCE          var4: [var5]]
          BENF             var6: [var7]
          [MODAL           var8]
          [TOPIC           var9]
          [CONTEXT         var10]
          {[modulators]}

var1   =   human_being_or_social_body
var3   =   service_
var4   =   human_being_or_social_body
var6   =   human_being_or_social_body
var8   =   process_, sector_specific_activity
var9   =   sortal_concept
var10  =   situation_
var2, var5, var7  =  geographical_location

c1)   MOVE   SUBJ     BRITISH_TELECOM
             OBJ      payg_internet_service
             BENF     (SPECIF customer_ BRITISH_TELECOM)
             date-1:  after-1-september-1998
             date-2:
```

When creating a *predicative occurrence (an instance of a template)* like c1 in the lower part of table 1, the role fillers in this occurrence *must conform to the constraints of the father-template*. For example, in occurrence c1, BRITISH_TELECOM is an individual, instance of the HClass concept company_: this last is, in turn, a specialization of human_being_or_social_body. payg_internet_service is a specialization of service_, a specific term of social_activity, etc. The meaning of the expression "BENF (SPECIF customer_ BRITISH_TELECOM)" in c1 is: the beneficiaries (role BENF) of the service are the customers of – SPECIF(ication) – British Telecom. The "attributive operator", SPECIF(ication), is one of the four operators used for the set up of the *structured arguments (expansions)* of conceptual predicates like MOVE, see Zarri (2009b: 68-70). In the occurrences, the two operators date-1 and date-2 materialize *the temporal interval normally associated with an elementary event*.

More than 150 templates are permanently inserted into HTemp; HTemp, *the NKRL ontology of events*, corresponds then to a sort of *'catalogue' of narrative formal structures*, which are easy to extend and customize. To supply now an at least intuitive idea of how a *complete narrative* is represented in NKRL, and returning to the table 1 example, let us suppose we would now state that: "We can note that, on March 2008, British Telecom *plans to offer* to its customers, in autumn 1998, a pay-as-you-go (payg) Internet service…", where the specific elementary event corresponding to the offer is still represented by occurrence c1 in table 1.

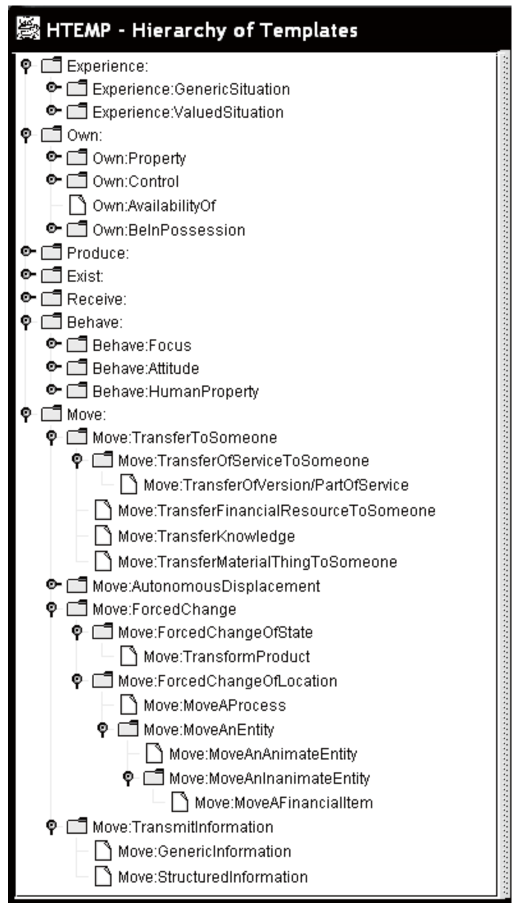

figure 2. 'MOVE' etc. branch of the HTemp hierarchy.

To encode correctly the new information, we must introduce first an *additional predicative occurrence* labeled as c2, see table 2, meaning that: "at the specific date associated with c2 (March 1998), it can be noticed, modulator obs(erve), that British Telecom *is planning* to act in some way" – the presence of a *second surface predicate* ("to plan") in the NL expression of the narrative denotes the presence of a second elementary event.

obs(erve) is a "*temporal modulator*", see Zarri (2009b: 71-72), used to identify a *particular timestamp* within the temporal interval of validity of an elementary event. We will then add a *binding occurrence* c3 labeled with a GOAL *Bn* operator, see the previous Section, and used to link together the conceptual labels c2 (the planning activity) and c1 (the intended result). The global meaning of c3 can be verbalized as: "The activity described in c2 is focalized towards (GOAL) the realization of c1". In agreement with the remarks at the end of the previous Section, c3 – *the representation of the global narrative* – can be represented *under tree form*, having GOAL as top node, and two branches with leaves L_1 = c2 and L_2 = c1.

table 2. Binding and predicative occurrences.

```
c2)   BEHAVE   SUBJ     BRITISH_TELECOM
               MODAL    planning_
                        { obs }
                        date1:  march-1998
                        date2:

Behave:ActExplicitly (1.12)

*c1)  MOVE   SUBJ     BRITISH_TELECOM
             OBJ      payg_internet_service
             BENF     (SPECIF customer_
                                  BRITISH_TELECOM)
             date-1:  after-1-september-1998
             date-2:

Move:TransferOfServiceToSomeone (4.11)

c3)   (GOAL  c2  c1)
```

The Querying/Inference Aspects

Reasoning in NKRL ranges from the *direct questioning* of a knowledge base of narratives represented in NKRL format – by means of *search patterns* p_i (formal queries) that unify information in the base thanks to the use of a *Filtering Unification Module* (*Fum*), see Zarri (2009: 183-201) – to *high-level inference procedures*. Thanks to the use of a powerful (and modular) *InferenceEngine*, these last make use of the richness of the representation to, e.g., set up *new relationships* among the narrative items stored in the base; a detailed paper on this topic is Zarri (2005).

The NKRL rules are characterized by the following general properties:

- All the NKRL high-level inference rules can be conceived as *implications* (biconditionals) of the type:

  X iff Y_1 and Y_2 … and Y_n . (4)

- In Eq. 4, X corresponds either to a *predicative occurrence* c_j (formal NKRL representation of an elementary event) or to a *search pattern* p_i (formal NLRL representation of a simple query) and $Y_1 \ldots Y_n$ – the NKRL translation of the *reasoning steps* that make up the rule – correspond to *partially instantiated templates*. They include then, see the upper part of Table 1 above, *explicit variables* of the form var_i.
- According to the usual conventions of *logic/rule programming*, *InferenceEngine understands each implication as a procedure*. This reduces *problems* of the form X to a *succession of sub-problems* of the form Y_1 and … Y_n.
- Each Y_i is interpreted in turn as a *procedure call* that tries to convert – using, in case, *backtracking procedures* – Y_i *into (at least) a successful search pattern* p_i. This last must then be able to unify (using the standard *Fum* module) one or several of the occurrences c_j of the NKRL knowledge base.
- The *success* of the unification operations of the pattern p_i derived from Y_i *means that the 'reasoning step' represented by* Y_i *has been validated*. *InferenceEngine* continues then its work *trying to validate the reasoning step corresponding to the sub-problem* Y_{i+1}.
- In line with the presence of the operator "and" in Eq. 4, the implication represented by Eq. 4 *is fully validated iff all the reasoning steps* $Y_1, Y_2 \ldots Y_n$ *are validated*.

As it appears from the above description, all the unification operations p_i/c_j required from the inference procedures make use only of the unification functions supplied by the Filtering Unification Module (*Fum*) introduced above. Apart from being used for the *direct questioning* operations, *Fum* constitutes as well, therefore, the *inner core* of the different *InferenceEngine* versions used in an NKRL framework. Details on the functioning of *InferenceEngine* – concerning, e.g., the way of avoiding infinite loops when working in a backward-chaining context – are given in Zarri (2005, 2009a).

From a practical point of view, the NKRL high-level inference procedures concern *mainly* two classes of rules, "transformations" and "hypotheses", see Zarri (2005).

Let us consider, e.g., the transformations. These rules try to '*adapt*', from a *semantic* point of view, a search pattern p_i that failed (that was unable to find a unification within the knowledge base) to the *real contents* of this base making use of a sort of *analogical reasoning*. They attempt then to *automatically 'transform'* p_i into one or more *different* $p_1, p_2 \ldots p_n$ that *are not strictly 'equivalent' but only 'semantically close'* (analogical reasoning) to the original one. In a transformation context, the *head* X of Eq. 4 is then represented by a search pattern, p_i.

Operationally, a transformation rule can be conceived as made up of a *left-hand side*, the "*antecedent*" – i.e. the formulation, in search pattern format, of the query to be transformed – and of one or more *right-hand sides*, the "*consequent(s)* – the NKRL representation(s) of one or more queries (search patterns) to be substituted for the given one. Denoting with A the antecedent and with Cs all the possible consequents, the transformation rules can then be expressed as:

$$A(var_i) \Rightarrow Cs(var_j), \quad var_i \subseteq var_j \qquad (5)$$

With respect to Eq. 4 above, X coincides now with A – a *search pattern* – while the reasoning steps $Y_1, Y_2 \ldots Y_n$ are used to produce the *search pattern(s)* Cs *to be used in place of* A. The restriction $var_i \subseteq var_j$ – all the variables declared in the antecedent A *must also appear* in Cs – assures the logical congruence of the rules. More formal details are given, e.g., in Zarri (2009: 212-216).

Let us consider a concrete example, which concerns a recent NKRL application about the 'intelligent' management of "storyboards" in the oil/gas industry. We want then ask whether, in a knowledge base where are stored all the possible *elementary and complex events* (narratives) related to the activation of a gas turbine, we can retrieve the information that a given oil extractor is running. In the absence of a direct answer we can reply by supplying, thanks to a transformation rule like that ($t11$) of table 3, other related events stored in the knowledge base, e.g., an information stating that the site leader has heard the working noise of the oil extractor, see figure 3. Expressed in natural language, this result can be paraphrased as: "The system cannot assert that the oil extractor is running, but it can certify that the site leader has heard the working noise of this extractor".

table 3. An example of 'transformation' rule.

$t11$: *"working noise/condition" transformation*

antecedent:

OWN	SUBJ	*var1*
	OBJ	property_
	TOPIC	running_

var1 = consumer_electronics, hardware_, surgical_tool, diagnostic_tool/system, small_portable_equipment, technical/industrial_tool

first consequent schema (*conseq1*):

EXPERIENCE	SUBJ	*var2*
	OBJ	evidence_
	TOPIC	(SPECIF *var3 var1*)

var2 = individual_person
var3 = working_noise, working_condition

second consequent schema (*conseq2*):

BEHAVE	SUBJ	*var2*
	MODAL	industrial_site_operator

Being unable to demonstrate directly that a given industrial apparatus is normally running, the fact that an operator can hear its working noise or note its operational aspect can be a proof of its running status.

With respect now to the *hypothesis rules*, these allow us to build up automatically a sort of *causal explanation* for an elementary event. In a hypothesis context, the head $X$of Eq. 4 is then represented by a predicative occurrence, c_j. Accordingly, the reasoning steps Y_i of Eq. 4 – called "condition schemata" in a hypothesis context – *must all be satisfied* (for each of them, at least one of the corresponding search patterns p_i must find a successful unification with the predicative occurrences of the base) *in order that the set of* $c_1, c_2 \ldots c_n$ *predicative occurrences retrieved in this way can be interpreted as a context/causal explanation of the original occurrence* c_j.

For example, to mention a well-known NKRL example, see Zarri (2005), let us suppose we have directly retrieved, in a querying-answering mode, information like: "Pharmacopeia, an USA biotechnology company, has received 64,000,000 dollars from the German company Schering in connection with an R&D activity" that corresponds then to c_j. We can then be able to automatically construct, using a hypothesis rule, a sort of causal explanation of this event by retrieving in the knowledge base information like: i) "Pharmacopeia and Schering have signed an agreement concerning the production by Pharmacopeia of a new compound" (c_1) and ii) "in the framework of this agreement, Pharmacopeia has actually produced the new compound" (c_2).

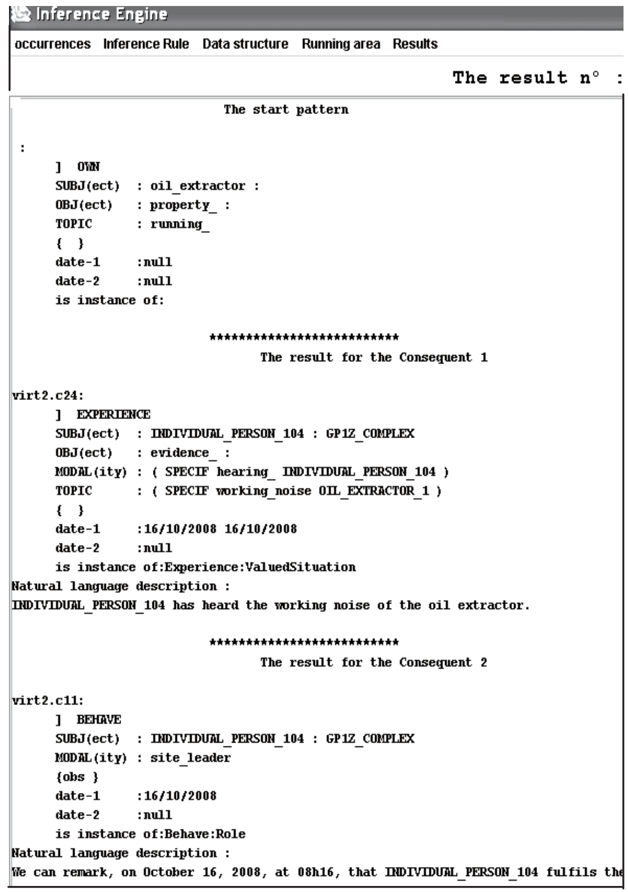

figure 3. Using InferenceEngine *in a 'transformation' context.*

An interesting, recent development of NKRL concerns the possibility of making use of the two above modalities of inference *in an 'integrated' way*, see Zarri (2009: 216-234). More exactly, *it is possible to make use of transformations when* InferenceEngine *is working within the hypothesis inference environment*. This means that, whenever a search pattern p_i is derived from a *condition schema of a hypothesis* to implement one of the steps of the reasoning process, we can use it '*as it is*' – i.e., in conformity with its father condition schema as this last has been coded when the inference rule has been built up – but also in *a 'transformed' form if the appropriate transformation rules exist within the system*.

The advantages are essentially of two types:

- From an utilitarian point of view, a hypothesis that was deemed to fail because of the impossibility of deriving successful p_i from one of its condition schemata *can now continue* if a transformed p_i is able to find an unification within the knowledge base, getting then new values for the hypothesis variables.
- From a more general point of view, this strategy allows us to explore in a systematic ways all the possible *implicit* relationships among the data in the base. A modality of the integrated strategy let us, in fact, *to transform all the* p_i *derived from the condition schemata of a hypothesis also in case of their successful unification with information in the base*. This permits, e.g., to confirm in many different ways the existence of relationships between people/entities.

The informal example of table 4 refers again to the oil/gas industry application. With this hypothesis, we should want to explain, see the *premise*, why an operator has activated a (particularly disturbing) "piping segment isolation procedure" in the context of, e.g., a gas leakage. The explication proposed is that i) a previous 'milder' procedure of the maintenance type has been executed (*cond1*), but this was unsuccessful (*cond2*); ii) the accident is a very serious one (*cond5*). In the absence of a predicative occurrence relating exactly this last fact, the p_i derived from *cond5 can be transformed* to obtain a confirmation of the gravity of the accident under several forms in the style of, e.g. "The gas leakage has a gas cloud shape" or "An alarm situation has been validated (*conseq1*) and the level of this alarm is 30% LEL, Low Explosion Level (*conseq2*)", etc.

table 4. 'Gas/oil' hypothesis in the presence of transformations.

(***premise***) An individual has carried out an "isolation" procedure in the context of an industrial accident.
(***cond1***) A different individual had carried out previously a (milder) "corrective maintenance" procedure.
(***cond2***) This second individual has experienced a failure in this corrective maintenance context.
(***cond3***) The first individual was a control room operator.
(***cond4***) The second individual was a field operator.
(***cond5***) The industrial accident is considered as a serious one.
- (**Rule t6, Consequent**) *The leakage has a gas cloud shape …*
- (**Rule t8, Consequent**) *A growth of the risk level has been discovered …*
- (**Rule t9, Consequent1**) *An alarm situation has been validated,* ***and***
- (**Rule t8, Consequent2**) *the level of this alarm is 30% LEL.*

Conclusion

In this paper we have introduced, first, some pragmatic criteria for, e.g., recognizing highly ambiguous entities like those elementary events that are the basic building blocks of a narrative. We have then introduced the main characteristics of a language, NKRL (Narrative Knowledge Representation Language), expressly specified for dealing with narrative and temporal information.

NKRL is a fully operational environment, developed thanks to several EC-funded projects. Successful applications in many different domains (from terrorism to the management of storyboards for the gas/oil industry…) have proved the practical utility of this tool.

References

Bal, M. 1997. *Narratology: Introduction to the Theory of Narrative* (2d ed). Toronto: University of Toronto Press.

Davidson, D. 1967. The Logical Form of Action Sentences. In: *The Logic of Decision and Action*. Pittsburgh (PA): University Press.

Higginbotham, J. 2000. On Events in Linguistic Semantics. In: *Speaking of Events*. Oxford: University Press.

Hoekstra, R., Liem, J., Bredeweg, B., and Breuker, J. 2006. Requirements for Representing Situations. In: *Proceedings of the OWLED*06 Workshop on OWL: Experiences and Directions* (vol. 216). Aachen: CEUR Publications.

Jackendoff, R. 1990. *Semantic Structures*. Cambridge (MA) : MIT Press.

Johnson, C.R., Fillmore, C.J., Wood, E.J., Ruppenhofer, J. Urban M., Petruck, M.R.L., and Baker, C.F. 2001. *The FrameNet Project: Tools for Lexicon Building* (Version 0.7). http://ccl.pku.edu.cn/doubtfire/semantics/FrameNet/theory/FrameNet_book.pdf

Kim, J. 1996. Events as Property Exemplifications. In: *Events*. Aldershot: Dartmouth Publishing.

Kipper, K., Korhonen, A., Ryant, N., and Palmer, M. 2008. A Large-Scale Classification of English Verbs. *Language Resources and Evaluation Journal* **42**, 21–40.

Liu, W., Liu, Z., Fu, J., Hu, R., and Zhong, Z. 2010. Extending OWL for Modeling Event-oriented Ontology. In: *Proceedings of the 2010 International Conference on Complex, Intelligent and Software Intensive Systems*. Los Alamitos (CA): IEEE Computer Society Press.

Mani, I., and Pustejovsky, J. 2004. Temporal Discourse Models for Narrative Structure. In: *Proceedings of the ACL Workshop on Discourse Annotation*. East Stroudsburg (PA): ACL.

Mizoguchi R., Sunagawa E., Kozaki K., and Kitamura Y. 2007. A Model of Roles within an Ontology Development Tool: Hozo. *Journal of Applied Ontology* **2**, 159-179.

Parson, T. 1990. *Events in the Semantics of English*. Cambridge (MA): The MIT Press.

Salguero, A.G., Delgado, C., Araque, F. 2009. Easing the Definition of N-Ary Relations for Supporting Spatio-Temporal Models in OWL. In: *EUROCAST 2009* (LNCS 5717). Berlin: Springer.

Zarri, G.P. 2005. Integrating the Two Main Inference Modes of NKRL, Transformations and Hypotheses. *Journal on Data Semantics (JoDS)* **4**, 304-340.

Zarri, G.P. 2009a. Using Rules in the Narrative Knowledge Representation Language (NKRL) Environment. In: *Handbook of Research on Emerging Rule-Based Languages and Technologies: Open Solutions and Approaches* (vol. 1). Hershey (PA): Information Science Reference.

Zarri, G.P. 2009b. *Representation and Management of Narrative Information, Theoretical Principles and Implementation*. London: Springer.