Proceedings of the Twenty-Eighth International Florida Artificial Intelligence Research Society Conference

Semantic/Conceptual Annotation Techniques Making Use of the Narrative Knowledge Representation Language (NKRL)

Gian Piero Zarri
Sorbonne University, STIH Laboratory – 75005 Paris, France
zarri@noos.fr

Abstract

This paper suggests the use of NKRL (Narrative Knowledge representation Language) as a tool for producing advanced "annotations" of textual and multimedia documents, i.e., annotations able to denote, in the best possible way, the global content (the 'meaning') of the documents.

Introduction

Labeling/annotation of textual and multimedia documents can be realized according to several (and quite different) approaches. The methodologies used in this domain are normally grouped into three rough categories: free form, linguistics-driven, and semantic/conceptual annotations.

Annotations in free form have been popularized by the Annotea project (Kahan et al. 2002). In Annotea, annotations are conceived as the means of introducing *generic remarks* (usually in natural language) about an existing document. Typically, free form annotations are: i) not formalized and impossible to exploit beyond simple keyword searches; ii) 'ephemeral'; iii) not very 'expressive' from a content representation point of view.

The Linguistic-Driven Annotations sub-domain is now quite popular thanks also to the success of GATE (General Architecture for Text Engineering), a tool developed at the Sheffield University (Cunningham 2002). In GATE, Computational Linguistics techniques are used to recognize and display the morphological, syntactic and semantic categories ("named entities" etc.) of specific terms or groups of terms that appear in the Natural Language (NL) documents examined. GATE's information extraction system (ANNIE) consists, e.g., of a set of modules including a tokenizer for lexical analysis, a gazetteer (dictionary), a sentence splitter, a part of speech tagger, a named entities transducer and a co-reference tagger. A sophisticated 'linguistic' annotation system implemented in the context of the EC-supported PARMENIDES project (Black et al. 2004) is organized in three *ontology-structured* layers: i) *structural annotations*, used to define the physical structure of the documents; ii) *lexical annotations*, used to mark 'interesting' text units like Named Entities, Temporal Expressions etc.; iii) *semantic annotations*, used to express the relationships that exist among lexical entities (e.g., lexically identified people can be associated with their organization and job title). The use of ontologies is quite common now in the linguistics-driven annotation domain see, e.g., the TELIX project (Rubiera et al. 2012) or the most recent versions of GATE (Cunningham et al. 2011).

The main function of the linguistics-driven annotation systems is the production of a *linguistically motivated, surface analysis* of the (*NL*) documents examined. This analysis is then *strongly dependent* from 'local' parameters like the particular language used, the lexical idiosyncrasies, the syntactic preferences (e.g., the active/passive alternation), the order of words in sentences, etc. The linguistic-driven annotations are not really concerned, then, with the representation of the "*deep meaning*" of textual/multimedia documents. This task is entrusted (at least in principle) to the *semantic/conceptual annotation tools*.

The use of *free keywords* and of *specialized thesauri* can be considered as the 'low level' of the semantic/conceptual annotation techniques. An evolution of this basic approach consists in passing from a simple "*keyword*" to a "*metadata*" approach. Metadata can be defined as *structured and machine understandable collections of lexical items to be used for expressing assertions about the organization and the contents of some sets of digital or non-digital documents*. There exists a wide variety of metadata formats for digital resources. They range from relatively simple ones, as the Dublin Core basic scheme, to more detailed formats like the Text Encoding Initiative (TEI, Wittern et al. 2009) and MARC (MAchine-Readable Cataloguing) (Fritz and Fritz 2003). Dublin Core has recently evolved into the Dublin Core Metadata Initiative (DCMI). The "*DCMI Metadata Terms*" (DCMI 2012) is a large set of metadata elements that includes the fifteen ones of the original Dublin Core along with several related properties and classes. Following the publication of the *RDF-compatible DCMI Abstract Model* (Nilsson et al.

2008), Dublin Core represents now one of the vocabularies most commonly used in a W3C/SW/RDF context.

One of the pioneering works about the use of ontologies for setting up *full conceptual annotations* of cultural heritage multimedia documents is described in Wielinga et al. (2001). More recent, ontologically-oriented work in the same domain concerns, e.g., the MuseumFinland and CultureSampo initiatives, developed in the context of the Finnish National SemanticWeb Ontology project. CultureSampo (Mäkelä et al. 2012) deals, e.g., with difficult problems like the management of multiple ontologies, multiple metadata schemas and cross-domain contents. In Arndt et al. (2009), the Authors describe COMM, an OWL DL multimedia ontology based on the MPEG-7 categories and concretely implemented as a composition of several "*multimedia patterns*". These denote small autonomous fragments of ontologies in the style of DOLCE (Masolo et al. 2003) that can be composed, specialized and reutilized.

The semantic/conceptual annotation techniques – those in particular that are *strongly ontology-based*, see Zarri (2014) for additional details – represent an important step towards the set-up of expressive annotations able to reflect, in some way, the *deep content (the 'meaning')* of the annotated documents. However, standard ontologies have difficulties when they have to deal with *real world, dynamic 'events'* ('actions', 'facts', situations', 'scenarios' etc.) that concern the concrete or intended behavior of some 'actors'. In such cases, the simple use of standard concepts and ontologies is not enough, and must be integrated by the description of the *mutual, dynamic relationships between the actors* – including then the description of their '*role*' in the context of the global actions, events etc.

In this paper, we argue then that the '*expressiveness*' of the present 'semantic/conceptual' annotation techniques could be greatly improved by adding, as in NKRL, the *Narrative Knowledge Representation Language* (Zarri 2009), an '*ontology of events*' to the traditional ontology of 'concepts' – and, more in general, by adding all those NKRL tools that are needed for dealing in full with the really 'pervasive' (Finlayson et al. 2010) "*narrative*" domain. Some basic information about NKRL will be supplied in the following Sections. A first Section will introduce NKRL's approach to the formalization of those "*elementary events*" that represent the building blocks of the notion of narrative. A narrative is usually interpreted, in fact, as *a sequence of logically structured and temporally and spatially bounded "elementary events" (a non-linear 'stream' of elementary events)* (Bal 1997). We will then supply some details about the implementation principles of NKRL; a final Section will represent a short "Conclusion".

Formalizing elementary events

The (*n*-ary) model used in NKRL to represent the elementary events – i.e., the core elements of any "narrative" (scenario, complex situation…) – can be denoted as:

$$(L_i\ (P_j\ (R_1\ a_1)\ (R_2\ a_2)\ \ldots\ (R_n\ a_n)))\ . \qquad (1)$$

In Eq. 1, L_i is the "*symbolic label*" identifying ('reifying') the formal description of a specific elementary event, P_j is a "*conceptual predicate*", R_k is a generic "*functional role*" and a_k the corresponding "*predicate arguments*". In the representation of an event like "Bill gives a book to Mary", the predicate P_j (of the MOVE type) will then introduce its *three* arguments a_k (JOHN_, MARY_ and BOOK_1) through *three functional relationships* (R_k roles) as SUBJ(ect), BEN(e)F(iciary) and OBJ(ect), the whole *n*-ary construction being reified through the symbolic label L_i.

An important property of NKRL concerns the fact that both the (unique) conceptual predicate of Eq. 1 and the associated functional roles are "*primitives*". Predicates P_j pertain then 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} – see Zarri (2011a) for the opposition "functional/semantic roles". The NKRL images of *specific elementary events* – obtained by producing *concrete instantiations* ("*predicative occurrences*") of structures in the style of Eq. 1 – are then, *at least partly*, a sort of *canonical representations.*

Several predicative occurrences – denoted by their symbolic labels L_i – can be assembled within the scope of second order structures called "*binding occurrences*", i.e., labeled lists made of a "binding operator *Bn*" with its arguments. The *Bn* operators – like ALTERN(ative), COORD(ination), CAUSE, GOAL, COND(ition) etc., see Zarri (2009: 91-98) – are used for *assembling the formal representations of elementary events denoted by their L_i labels into larger structures* (narratives, scenarios etc.). The general expression of a binding occurrence is then:

$$(Bn_k\ arg_1\ arg_2\ \ldots\ arg_n)\ . \qquad (2)$$

The arguments arg_i can correspond directly to L_i labels or recursively denote sets of labeled lists in Eq. 2 format.

Some implementation details

The knowledge representation aspects

NKRL adds to the usual "*ontology of concepts*" – called HClass (hierarchy of classes) in NKRL's terms, and used to define the a_k terms of Eq. 1 above – an "*ontology of elementary events*". This last is *a new type of hierarchical organization – called HTemp, hierarchy of templates – where the nodes consist of* n-*ary structures built around the basic formal core represented by Eq. 1*. Templates can be conceived as the *canonical representation of generic classes of elementary events* like "move a physical object", "be present in a place", "send/receive a message", etc. More than 150 templates are permanently inserted into HTemp, which corresponds then to a sort of *catalogue* of narrative formal structures, easy to extend and customize.

When a *specific elementary event* must be represented, the corresponding *predicative occurrence* is derived by instantiating the appropriate HTemp template. To annotate then a news story fragment (an elementary event) like:

"British Telecom will offer its customers a pay-as-you-go (payg) Internet service in autumn 1998", we must select firstly in HTemp the template corresponding to "supply a service to someone", see the upper part of Table 1.

As it can be seen from this table in the complete, formal representation of the templates, the arguments of the predicate (the a_k terms in Eq. 1) are represented *by variables with associated constraints*. These last *are expressed in turn as HClass concepts or combinations of concepts*: the two ontologies, HTemp and HClass *are strictly intermingled*. In a *predicative occurrence* like c1 in Table 1, the role fillers *must conform to the constraints of their father-template*. For example, BRITISH_TELECOM is an individual, instance of the HClass concept company_ that is, in turn, a specialization of human_being_or_social_body, etc. The meaning of the expression "BENF (SPECIF customer_ BRITISH_TELECOM)" is: the beneficiaries of the service are the customers of – SPECIF(ication) – British Telecom. The two operators date-1 and date-2 materialize *the temporal interval normally linked to an elementary event*, see Zarri (2009: 76-86) in this context.

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 Somone"

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:	

To supply an at least intuitive idea of how a *complete narrative* is represented in NKRL, let us suppose we would now *annotate* a more complete version of the previous news: "We can note that, on March 1998, 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. We must then introduce first an *additional predicative occurrence* labeled as c2, and meaning that: "at the specific date associated with c2 (March 1998), it can be noticed that British Telecom *is planning* to act in some way". We will then add a *binding occurrence* c3 labeled with GOAL (a *Bn* operator) and used to link together the labels c2 (the planning activity) and c1 (the intended result): "c3 (GOAL c2 c1)". The meaning of c3 can be verbalized as: "The activity described in c2 is focalized towards (GOAL) the realization of c1".

The querying/inference aspects

Reasoning in NKRL ranges from the *direct questioning* of a knowledge base (*KB*) of annotations in NKRL format to *high-level inference procedures*. Direct questioning is implemented by means of *search patterns* p_i (corresponding to *specialized/partially instantiated templates*) that unify, in a subsumption-oriented style, information in the base thanks to the use of a *Filtering Unification Module*, *Fum* (Zarri 2009: 183-201). High-level inference procedures utilize a backward-chaining *InferenceEngine* that makes use of the richness of the representation to, e.g., set up *new relationships* among the items stored in the base.

'Standard' search patterns p_i are also *automatically* generated by *InferenceEngine* as the final forms of the different *reasoning steps* that make up the high-level inference procedures. These last are represented *mainly* by "transformations" and "hypotheses" (Zarri 2009: 201-239).

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 *KB*) to the *real contents* of this base making use of *analogical reasoning*. The rules attempt then to *automatically 'transform'* p_i into one or more *different* p_1, p_2 ... p_n that *are not strictly 'equivalent' but only 'semantically close'* to the original one. A transformation rule is made 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 representation of the pattern(s) to be substituted for the given one. Denoting with *A* the antecedent and with Cs_i the possible consequents, these rules can be expressed as:

$$A(var_i) \;\Rightarrow\; Cs_i(var_j), \quad var_i \subseteq var_j \qquad (3)$$

We can now see a concrete example, which concerns a recent application about the 'intelligent' management of NKRL annotations of "storyboards" in the oil/gas industry, see Zarri (2011b). We want to ask whether, in a *KB* where are stored some events related to the activation of a gas turbine, we can retrieve the information that a given oil extractor is running. In the absence of direct answer we can reply by supplying, thanks to a rule like t11 (Table 2), other related events, e.g., *information stating that the site leader has heard the working noise of the extractor*. 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 its working noise".

The *hypothesis rules* allow us to build up automatically a sort of '*causal explanation*' for an event (predicative occurrence) retrieved within a NKRL knowledge base. These rules can be expressed as *biconditionals* of the type:

X iff Y_1 and Y_2 ... and Y_n , (4)

where the 'head' X of the rule corresponds to a predicative occurrence c_j to be 'explained' and the 'reasoning steps' Y_i – called "condition schemata" in a hypothesis context – *must all be satisfied.* This means that, *for each of them*, at least one '*successful*' search patterns p_i must be (automatically) derived by *InferenceEngine* in order to find, using *Fum*, a *successful unification* with some information of the base. *In this case, the set of* c_1, c_2 ... c_n *predicative occurrences retrieved by the condition schemata* Y_i *thanks to their conversion into* p_i *can be interpreted as a context/causal explanation of the original occurrence* c_j (X).

Table 2. An example of transformation rule.

t11: *"working noise/condition" transformation*

antecedent:

OWN	SUBJ	*var1*
	OBJ	property_
	TOPIC	running_

var1 = consumer_electronics, hardware_, diagnostic_tool/system, small_portable_equipment, surgical_tool, 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 an industrial apparatus is running, the fact that an operator hears its working noise or notes its working aspect can be a proof of its running status.

To mention a well-known NKRL example, let us suppose we have directly retrieved, in a querying-answering mode, the information: "Pharmacopeia, a USA biotechnology company, has received 64,000,000 dollars from the German company Schering in connection with an R&D activity"; this information corresponds then to c_j (X). We can be able to automatically construct, using a "hypothesis" rule, a sort of 'causal explanation' for this event by retrieving in the *KB* information like: i) "Pharmacopeia and Schering have signed an agreement concerning the production by Pharmacopeia of a new compound", c_1 (Y_1) and ii) "in the framework of this agreement, Pharmacopeia has actually produced the new compound", c_2 (Y_2).

Conclusion

This paper focuses on the possible improvements to the "*expressiveness*" of the present 'semantic/conceptual' annotation techniques that could be obtained by using some of the high-level knowledge representation and querying/inference techniques proper to NKRL, the "Narrative Knowledge Representation Language". Some important NKRL's features – illustrated through the use of concrete examples – are described sketchily in this paper.

References

Arndt, R., Troncy, R., Staab, S., and Hardman, L. 2009. COMM: A Core Ontology for Multimedia Annotation. In: *Handbook on Ontologies* (2nd ed.), 403-421. Berlin: Springer.

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

Black, W.J., Jowett, S., Mavroudakis, T., McNaught, J., Theodoulidis, B., Vasilakopoulos, A., Zarri, G.P., and Zervanou, K. 2004. Ontology-Enablement of a System for Semantic Annotation of Digital Documents. In: *Proc. of SemAnnot 2004*. Trier: CEUR (http://ceur-ws.org/Vol-184/semAnnot04-05.pdf, accessed 10/2/2015).

Cunningham, H. 2002. GATE, A General Architecture for Text Engineering. *Computers and the Humanities* **36**: 223-254.

Cunningham, H., Maynard, D., Bontcheva, K., Tablan, V., Aswani, N., Roberts, I., Gorrell, G., Funk, A., Roberts, A., Damljanovic, D., Heitz, T., Greenwood, M., Saggion, H., Petrak, J., Li, Y., and Peters, W. (2011). *Text Processing with GATE (Version 6)*. Sheffield: Department of Computer Science of the University.

DCMI. 2012. *DCMI Metadata Terms*. Dublin (OH): OCLC Research (http://dublincore.org/documents/dcmi-terms/, accessed 10/2/2015).

Finlayson, M.A., Gervás, P., Mueller, E., Narayanan, S., and Winston, P., eds. 2010. *Computational Models of Narratives – AAAI 2010 Fall Symposium (TR FS-10-04)*. Menlo Park (CA): AAAI Press.

Fritz, D.A., and Fritz R.J. 2003. *MARC 21 for Everyone: A Practical Guide*. Chicago: American Library Association.

Kahan, J., Koivunen, M.R., Prud'hommeaux, E., and Swick, R.R. 2002. Annotea: An Open RDF Infrastructure for Shared Web Annotations. *Computer Networks* **39**: 589-608.

Mäkelä, E., Hyvönen, E., and Ruotsalo, T. 2012. How to Deal with Massively Heterogeneous Cultural Heritage Data, Lessons Learned in CultureSampo. *Semantic Web – Interoperability, Usability, Applicability* **3**: 85-109.

Masolo, C., Borgo, S., Gangemi, A., Guarino, N., and Oltramari, A. 2003. *WonderWeb D18*. Trento: Laboratory for Applied Ontology (http://www.loa.istc.cnr.it/old/Papers/D18.pdf, accessed 10/2/2015).

Nilsson, M., Powell, A., Johnston, P., and Naeve, A. 2008. *Expressing Dublin Core Metadata Using RDF*. Dublin (OH): OCLC Research. (http://dublincore.org/documents/dc-rdf/, accessed 10/2/2015).

Rubiera, E., Polo, L., Berrueta, D., and El Ghali, A. 2012. TELIX: An RDF-Based Model for Linguistic Annotation. In: *The Semantic Web – Proc. of ESWC 2012* (LNCS 7295), 195-209. Berlin: Springer.

Wielinga, B.J., Schreiber, A.Th., Wielemaker, J., and Sanderberg, J.A.C. 2001. From Thesaurus to Ontology. In: *Proc. of the 1st Int. Conference on Knowledge Capture (K-CAP '01)*, 194-201. New York: ACM Press.

Wittern, C., Ciula, A., and Tuohy, C. 2009. The Making of TEI P5. *Literary & Linguistic Computing* **24**: 281-296.

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

Zarri, G.P. 2011a. Differentiating Between "Functional" and "Semantic" Roles in a High-Level Conceptual Data Modeling Language. In: *Proc. of FLAIRS-24*. Menlo Park (CA): AAAI Press.

Zarri, G.P. 2011b. Knowledge Representation and Inference Techniques to Improve the Management of Gas and Oil Facilities. *Knowledge-Based Systems (KNOSYS)* **24**: 989–1003.

Zarri, G.P. 2013. Advanced Computational Reasoning Based on the NKRL Conceptual Model. *Expert Systems With Applications (ESWA)* **40**: 2872-2888.

Zarri, G.P.. 2014. Conceptual and Content-based Annotation of (Multimedia) Documents. *Multimedia Tools and Applications* **72**: 2359-2391.