Verb Interpretation for Basic Action Types: Annotation, Ontology Induction and Creation of Prototypical Scenes

Francesca Frontini, Irene De Felice, Fahad Khan, Irene Russo, Monica Monachini[1] Gloria Gagliardi, Alessandro Panunzi[2]

(1) ILC CNR Pisa, {francesca.frontini, irene.defelice, fahad.khan, irene.russo, monica.monachini}
@ilc.cnr.it
(2) University of Florence, gloria.gagliardi@unifi.it

ABSTRACT

In the last 20 years dictionaries and lexicographic resources such as WordNet have started to be enriched with multimodal content. Short videos depicting basic actions support the user's need (especially in second language acquisition) to fully understand the range of applicability of verbs. The IMAGACT project has among its results a repository of action verbs ontologically organised around prototypical action scenes in the form of both video recordings and 3D animations. The creation of the IMAGACT ontology, which consists in deriving action types from corpus instances of action verbs, intra and cross linguistically validating them and producing the prototypical scenes thereof, is the preliminary step for the creation of a resouce that users can browse by verb, learning how to match different action prototypes with the correct verbs in the target language. The mapping of IMAGACT types onto WordNet synsets allows for a mutual enrichment of both resources.

Interpretazione dei verbi per tipi azionali di base: annotazione, induzione di ontologia e creazione di scene prototipiche

Negli ultimi venti anni dizionari e risorse lessicografiche come WordNet sono stati arricchiti con contenuto multimediale. Brevi video in grado di rappresentare azioni di base supportano i bisogni degli utenti (in particular modo per quanto riguarda l' acquisizione della seconda lingua) nel comprendere l' ambito di applicabilità dei verbi. Il progetto IMAGACT ha tra i suoi risultati una base di dati di verbi d'azione ontologicamente organizzati e raffiguranti scene che riproducono azioni prototipiche sottoforma di registrazioni video e animazioni 3D. La creazione dell' ontologia IMAGACT che consiste nella derivazione di tipi azionali da istanze di verbi d'azione estratte da un corpus, nella loro validazione intra e crosslinguisticamente e nella conseguente produzione di scene prototipiche, è il passaggio preliminare per la creazione di una risorsa che gli utenti possono consultare partendo dal verbo, imparando come alllineare differenti prototipi d'azione con il verbo corretto nella lingua da apprendere. Il *mapping* dei tipi di IMAGACT sui *synsets* di WordNet consente un arricchimento reciproco di entrambe le risorse.

KEYWORDS : ontology of actions, lexical resource, 3D animations
KEYWORDS IN ITALIAN : ontologia di azioni, risorse lessicali, animazioni 3D

1   Introduction

In the last 20 years dictionaries and lexicographic resources such as WordNet have started to be enriched with multimodal content (e.g. pictorial illustrations, animations, videos, audio files). Pictures are effective in conveying the meaning of denotative words such as concrete nouns, while for abstract relations (instantiated by prepositional meanings) schematic illustrations can depict several semantic properties. Conveying the meaning of verbs with static representations is not possible; for such cases the use of animations and videos has been proposed (see Stein 1991 cited in Lew 2010). Short videos depicting basic actions support the user's need (especially in second language acquisition) to fully understand the range of applicability of verbs i.e. to start with a mental image of an action and from this image find out the L2 verb(s) that can be used to predicate that action. This process involves semantic and pragmatic comparisons that occur in the mind of the learner, with considerations respecting the type of movement involved, the instrument/tool that can be used, the duration, the strength of the movement etc.

In this paper we introduce the IMAGACT project and its results: a repository of action verbs ontologically organised around prototypical action scenes in the form of both video recordings and 3D animations. The focus of IMAGACT is on action verbs, because in all language modalities they bear basic information that should be processed in order to make sense of a sentence. Especially in speech, they are the most frequent structuring elements (Moneglia and Panunzi, 2007), but unfortunately no one-to-one correspondence can be established between an action verb, conceived as a lexical entry, and an action type, conceived as an ontological entity.

In order to bridge this gap 500 English and Italian action verbs have been analysed in their different contexts of use in corpora and grouped into action types according to their internal variation. Types representing the same prototypical actions are then gathered together under the same scene and represented in 3D animations, generated ad hoc which thus illustrate the different uses of action verbs across languages (see Figure 1).

For instance, the English verb *to roll* can refer to qualitatively different actions. In some uses the agent changes the form of the object (B and 1), in some other uses the agent moves himself in space (C and 2), and in other cases, the agent moves the object in space, applying a force to it (D and 3):

(1)   John rolls the poster into a tube.
(2)   John rolls onto his side.
(3)   John rolls the barrel.

In short, different action types occur in the above examples. This judgment is confirmed by the productivity of each action type. Despite the fact that the predicate is applied to different objects, humans are able to judge whether the same action is performed or not by reading a set of sentences:

(1a)   John rolls the poster / his sleeve/ the pants up.
(2a)   John /Mary / the horse rolls.
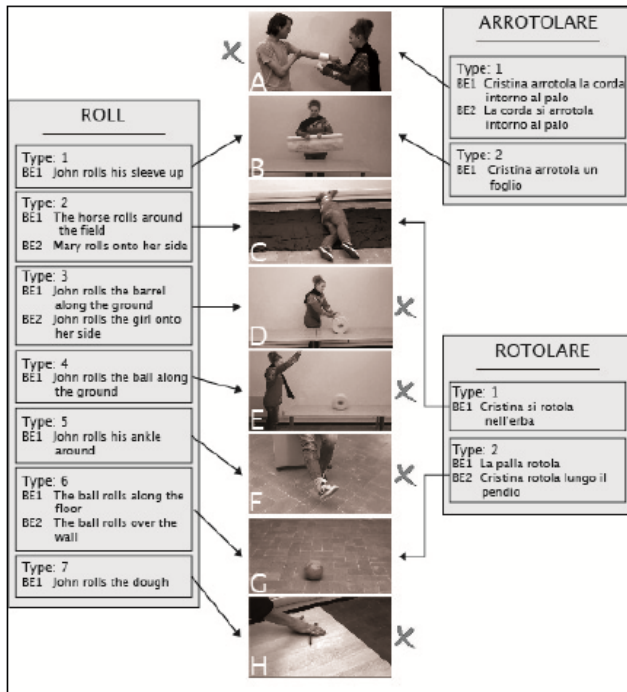(3a)   John rolls the barrel / the cylinder.

FIGURE 1 - Cross-linguistic gallery of scenes representing the variation of *to roll, arrotolare* and *rotolare*

In other words, *to roll* has several interpretations corresponding to the different action types, and none of these types can be considered more appropriate than the others in characterizing the meaning of the verb. Each one could be a prototypic instance of the verb (Givon, 1986).

We call general verbs all natural language action verbs that share this property. In the case of general verbs, ordinary language does not mirror the ontology of action and the lemma does not specify the referred ontological entity. As shown by Figure 1, the different types of a general verb may map onto different verbs in other languages. This causes huge problems for second language acquisition since each language categorises the space of action in its own way. Figure 1 is an example of the relation of English and Italian verbal entries with respect to the same *continuum*.

The targets of the IMAGACT resource are L2 learners of the supported languages (focus on Italian) who can browse the resource by verb, learning how to match different action prototypes with the correct verbs in the target language.

In the following paragraphs we shall describe the procedure for the creation of the IMAGACT ontology, which consists in deriving action types from corpus instances of action verbs, intra and cross linguistically validating them and producing the prototypical scenes thereof. Criteria

applied for the creation of prototypical scenes will also be investigated. Finally, the possibility of mapping of IMAGACT types onto WordNet synsets, thus allowing for a mutual enrichment of both resources. We will end with conclusions and ideas for future work.

2   Related Works

The importance of providing visual support for lexical and ontological resources is becoming more and more evident. Ontologies like SUMO[1] provide links to pictures from external sources (often Wikimedia) to add a visual illustration of many of its concepts. DBpedia also contains links to pictures, which are already part of the information derived from each Wikipedia entry. Image-net[2] goes even further, presenting itself as a veritable image database organised according to the WordNet hierarchy.

In traditional dictionaries words are explained with words, using a definition or an equivalent word (for bilingual dictionaries); definitions as paraphrases of lexical units through syntactic construction (with or without examples) are common also in lexical resources such as WordNet. In electronic dictionaries a wide usage of other means (such as pictorial illustrations, pictures, animations, videos, audio files) is possible and paves the way for multimodal lexicographic resources. If pictorial illustrations are effective for nouns (in particular for plants, animals and common objects), their utility for complex actions and the abstract or figurative meaning of words is less predictable. Adamska-Sałaciak (2008) (working on lexicography from a cognitive linguistics perspective) suggests that the inclusion of schematic graphs to represent the meaning of prepositions in dictionaries is useful. Animated illustrations are effective because they provide user-friendly representation of stages or the progression of an action and, together with videos, constitute the better modality for presenting verbal meanings, even if this is still an underinvestigated issue. Video sequences can convey information about situational contexts but are rather costly in terms of storage space and their realization is not easy (i.e. several semiotic principles should be followed for their realization).

Yet in all these resources entries are linguistically or conceptually motivated. Images are linked to concepts, synsets or lexical entries, which provide the hierarchical structure to the resource. None has, to our knowledge, attempted to do the inverse; that is to build a veritable visual ontology, where the types are visually represented, and semantic and lexical information is dependent to visual types. In the IMAGACT ontology each type is represented by a prototypical scene, specifically one produced with 3D animation techniques in order to describe in a salient way one prototypical action.

3   The IMAGACT project

The IMAGACT project uses both corpus-based and competence-based methodologies for simultaneous extraction of a language independent action inventory from spontaneous speech corpora of different languages.

The IMAGACT infrastructure faces key issues in ontology building. It grounds productive translation relations since it distinguishes the primary usage of verbs from their metaphorical or

---

[1] sigma.ontologyportal.org:4010/sigma/Browse.jsp?lang=EnglishLanguage&flang=SUO-KIF&kb=SUMO&term=Pump

[2] www.image-net.org

phraseological extensions; it allows easy identification of types in the variation, it is cross-linguistic in nature, it derives from the actual use of language but it can be freely extended to other languages through competence-based judgments and it is therefore suitable for filling gaps in lexical resources.

The IMAGACT database focuses on high frequency action verbs, which can provide sufficient variation in spoken corpora; i.e. roughly 500 verbs referring to actions which represent the full basic action oriented verbal lexicon. In order to maximize the probability of occurrence of relevant action types, IMAGACT identifies the variation of this set in parallel on two spoken corpora:

- a 2 million word English corpus, taken from the British National Corpus;

- a collection of spoken Italian corpora with 1.6 million words in total (LABLITA corpus, Cresti and Moneglia, 2005; LIP, De Mauro et al., 1993; CLIPS corpus).

3.1   The IMAGACT annotation framework

The annotation procedure is structured into two main steps, standardization & clustering of occurrences and types annotation & assessment, accomplished by annotators with the assistance of a supervisor. The first task is to examine and interpret verb occurrences in the oral context, which is frequently fragmented and may not provide enough semantic evidence for an immediate interpretation. To this end the infrastructure allows the annotator to read the larger context of the verbal occurrence in order to grasp the meaning (Figure 2 presents two of the occurrences of *to roll* in the corpus). The annotator represents the referred action with a simple sentence in a standard form for easy processing. This sentence must be in the positive form, in the third person, present tense, active voice and must fill the essential argument positions of the verb (possible specifiers that are useful in grasping the meaning are placed in square brackets). Basic level expressions (Rosch 1978) are preferred or otherwise a proper name is used and word order in sentences must be linear, with no embedding and/or distance relationships.

Crucially, along with the standardization, the annotator assigns each occurrence to a "variation class" thus determining whether or not it conveys the verb's meaning. This is what we mean by a PRIMARY occurrence. This task is accomplished through a synthetic judgment which exploits the semantic competence of the annotator (Cresswell 1978) and is given in conjunction with Wittgenstein's hypothesis on how word extensions can be learned (Wittgenstein 1953). The occurrence is judged PRIMARY according to two main operational criteria: a) it refers to a physical action; b) it can be presented to somebody who does not know the meaning of the verb V, by asserting that "the referred action and similar events are what we intend with V".
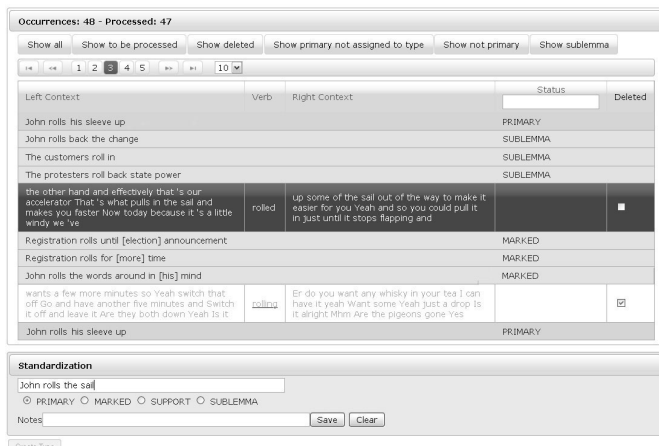
Occurrences: 48 - Processed: 47

Show all | Show to be processed | Show deleted | Show primary not assigned to type | Show not primary | Show sublemma

|◄ | ◄◄ | 1 2 **3** 4 5 | ►► | ►| | 10 ▾ |

Left Context	Verb	Right Context	Status	Deleted
John rolls his sleeve up			PRIMARY	
John rolls back the change			SUBLEMMA	
The customers roll in			SUBLEMMA	
The protesters roll back state power			SUBLEMMA	
the other hand and effectively that 's our accelerator That 's what pulls in the sail and makes you faster Now today because it 's a little windy we 've	rolled	up some of the sail out of the way to make it easier for you Yeah and so you could pull it in just until it stops flapping and		☐
Registration rolls until [election] announcement			MARKED	
Registration rolls for [more] time			MARKED	
John rolls the words around in [his] mind			MARKED	
wants a few more minutes so Yeah switch that off Go and have another five minutes and Switch it off and leave it Are they both down Yeah Is it	rolling	Er do you want any whisky in your tea I can have it yeah Want some Yeah just a drop Is it alright Mhm Are the pigeons gone Yes		☑
John rolls his sleeve up			PRIMARY	

Standardization

John rolls the sail

◉ PRIMARY ○ MARKED ○ SUPPORT ○ SUBLEMMA

Notes [   ]   Save | Clear

Create Type

FIGURE 2 - Verb occurrence and Standardization box

The occurrence is judged MARKED otherwise, as with "John rolls the words in his mind".

Only occurrences assigned to the PRIMARY variation class make up the set of Action Types stored in the ontology. To this end they must be clustered into families which constitute the productive variation of the verb predicate. The workflow thus requires the examination of the full set of standardized primary occurrences recorded in the corpus, whose meaning is now clear.

The infrastructure is designed to allow the annotator to create types ensuring both cognitive similarity among their events and pragmatic differences between them. The overall criterion for type creation is to keep granularity to its minimal level, assigning instances to the same type as long as they fit with one "best example". Clustered sentences should be similar as regards:

- the possibility to extend the occurrence by way of similarity with the virtual image provided by the best example (Cognitive Constraint);
- "equivalent verbs applied in their proper meaning" i.e. the synset (Fellbaum 1998) (Linguistic Constraints);
- involved Action schema.

Among the occurrences the annotator chooses the most representative as best examples of the recorded variation, creates types headed by one (or more) best example(s), and assigns each individual standardization to a type by dragging and dropping. The infrastructure assists the annotator in the task by showing the types that have been created so far and the equivalent verbs used to differentiate them.

The assigned instances can be shown by type and best example according to the annotator's needs. The infrastructure also provides functionality for making easy revisions to hypotheses (by showing instances not yet assigned, showing all instances, the verification of Marked variation, editing/merging/splitting types etc.).

The approach underlying the annotation strategy does not require a priori any inter-annotator agreement in this core task, which is strongly underdetermined, and rather relies on a supervised process of revision.

Once all occurrences have been processed, negotiation with a supervisor leads to a consensus on the minimal granularity of the action types extended by the verb in its corpus occurrences. The verification criteria are practical: the supervisor verifies for each type that it cannot be referred to as an instance of another without losing internal cohesion. The operational test checks if it is understandable that the native speaker is referring to the event by pointing to the prototype. The supervisor considers the pragmatic relevance of these judgments and keeps the granularity accordingly.

The relation to images of prototypical scenes provides a challenging question in restricting granularity to a minimal family resemblance set: "can you specify the action referred to by one type as something like the best example of another?". Granularity is kept when this is not reasonable.

Once types are verified the infrastructure presents the annotator with the "Types Annotation & Assessment" interface. Conversely, in this task the annotator assesses that all instances gathered within each type can indeed be extensions of its best example(s), thus validating its consistency. Those that aren't are assigned to other types.

FIGURE 3 - Types Annotation and Assessment

The assessment runs in parallel with the annotation of the main linguistic features of a type. More best examples can be added in order to represent all the thematic structures of a verb which satisfies that interpretation. As shown in Figure 3 the thematic grid must be filled, by writing each argument in a separate cell and selecting a role-label from the adjacent combo-box. The tag-

set for thematic role annotation is constituted by a restricted set of labels derived from current practices in computational lexicons. We are using Palmer's Tagset in VerbNet with adaptations.

Each best example is also annotated with an aspectual class which is assigned by means of the Imperfective Paradox Test (Dowty, 1979). Aspect can assume three values: event, process or state. Sentences that are judged peripheral instances of the type can be marked, thus identifying fuzziness in pragmatic boundaries. The annotation procedure ends when all proper occurrences of a verb have been assessed. The annotator produces a "script" for each type and delivers the verb annotation to the supervisor for cross-linguistic mapping.

3.1.1 Description of the methodology of interlinguistic validation

The direct representation of actions through scenes that can be interpreted independently of language allows the mapping of lexicons from different languages onto the same cross-linguistic ontology.

Working with data coming from more than one language corpus, IMAGACT must produce a language independent type inventory. For instance, in the case of *to roll* action types must be consistent with those extended by the Italian verb *rotolare/arrolotolare*, which in principle could be roughly equivalent. Therefore the supervisor will face two lists of types independently derived from corpora annotation. In this scenario, setting the cross-linguistic relations among verbal entries relies on the identification of a strict similarity between the Types that have been identified (and not through the active writing of a definition). The task is mapping similar types onto one prototypic scene that they can be an instance of.

Figure 1 roughly sketches the main types derived from the annotation of *to roll* and *rotolare / arrotolare* and their mapping onto scenes. The supervisor should recognize for instance, that type 2 of *to roll* and type 1 of *rotolare* are instances of the same prototype. The supervisor will accordingly produce a scene (scene C here). Cross-linguistic mapping allows us to predict relevant information which does not emerge from simple corpus annotation. For instance some types of *rotolare* may never occur in the English corpus, but native English speakers can recognize from the scene that they too are a possible extension of *to roll*. The mapping of the verb onto that type will therefore be established, providing competence based information. Mappings are not always possible: in this case the native speaker recognizes that T1 of *to roll* cannot be extended by *rotolare* while *arrotolare* is applicable. In other words the infrastructure and the methodology embodied in it allow the identification of the pragmatic universe of action and of how different languages parse it. This result is obtained in a Wittgenstein-like scenario without the comparison of definitions. The use of prototypic images bypasses this complex problem and permits the identification of the focal pragmatic variation of general verbs and their differentials in different languages.

Notice that this first mapping is performed on the basis of Types only. Its productivity must be then validated at the level of each single instance. A second step of interlinguistic validation consists in asking mother tongue informants what verb(s) should be applied in their language to each scene and whether the verb(s) is applicable to the set of English/Italian sentences headed by that scene.

Crucially, the informant will verify whether or not the choice is correct for all arguments retrieved from the corpus and assigned to that type and in doing so will verify to which extent the pragmatic concepts stored in the ontology are productive i.e. they permit generalizations at a

cross-linguistic level. This means that in IMAGACT a concept is valid for cross-linguistic reference to action if, independently of the language, the verb that is applied to the prototypic instance can also be applied to all sentences gathered in it.

The cross linguistic validation is performed in parallel on English and Italian sentences gathered within each entry and it generates a data set of parallel sentences. A competence based extension to other languages (Spanish and Chinese Mandarin) is also in progress, and consists in identifying a verb in the target language for each type of the source language and verifying the applicability to all instances in the target language, without actually producing sentences in the target language.

The interlinguistic validation of types is a very crucial phase of the IMAGACT project. Distinguishing families of usages of general verbs from the granular variations allows us to discover productive cross-linguistic relations, thus validating the ontology entries in the real world.

4   From words to videos: methodology

Once types of actions referred to by action verbs have been identified and the scripts have been produced for the best examples, with cross-linguistic equivalences established, the supervisor produces a prototypical scene.

Actors perform the action described in the script or an equivalent action. The scene is recorded according to the following requirements, which are intended to reduce ambiguity and to trigger the preferred interpretation:

- Use of real-world objects instead of abstract/generic forms
- Minimal, necessary background information
- The scene is produced as an uninterrupted shot ("long take")
- The action is performed with its usual temporal span (no slow-motion)
- The sequence is edited to focus on the sole relevant nucleus of the performed action (3-7 seconds)

The semiotic relevance of each scene and its capacity to elicit the appropriate verb is scrutinized by more than three experts before storage in the database.

Subsequently a 3D animation is created from the videos, in order to make the scene even less ambiguous. The animation software used for the production of 3D videos is Autodesk MAYA[3].

An animation must be equivalent to the real scene for what concerns its possible interpretation, but not necessarily equivalent with respect to the used objects.

5   Mapping IMAGACT onto WordNet

We are currently dealing with another task, that is to establish a link between IMAGACT and WordNet.

---

[3] The output format is H.264/mpeg-4, with framesize 1024*576.

WordNet is one of the best-known lexical resources and it contains one of the most complete verbal ontologies of any lexical resource, not only in terms of lexical entries, but also for the number of relations among verbs (hyponymy/hypernymy, troponymy, entailment). It is therefore very useful to investigate how IMAGACT maps onto WordNet. A mapping of both resources would lead to a reciprocal enrichment of several aspects: for instance IMAGACT does not show semantic relations among verbs, nor does it use definitions/glosses to define actions or action types, while WordNet does; on the other side WordNet does not distinguish between primary and marked senses, often confusing proper uses with metaphorical or idiomatical ones. Furthermore, WordNet defines horizontal relations among senses (synsets) with glosses, while IMAGACT uses scenes to represent the event type which different verbs can refer to in similar contexts (equivalent verb classes). So in case of perfect matching between an action type and a synset, IMAGACT videos would be enriched by WN glosses, and WN glosses could be more intuitively understood if visually represented.

It is also important to stress that WordNets have been now produced for many languages (and sometimes connected one to another: see for example EuroWordNet, GlobalWordNet projects). This would allow in the future the extension of the mapping to new languages, once they have been implemented in IMAGACT. Furthermore, we can imagine that if different WordNet ontologies are mapped onto the same IMAGACT interlinguistical ontology, they will be automatically linked one to another, and this will be of great benefit to the multilingual projects cited above.

As we said above the ontology of action types has already been completed by extracting data from Italian corpus annotation, therefore a first mapping of Italian action types onto ItalWordNet senses has been attempted. For every IMAGACT action verb, we compared the action types with the senses of the corresponding ItalWordNet lexical entry and with their related synsets.

We have already mapped about 150 Italian action types onto ItalWordNet. In some cases, especially when the verb refers to a very specific action (e.g. *stirare, to iron*) or it has a strong prototypical meaning (e.g. *camminare, to walk*), as often happens with activity verbs, the verb has only one IMAGACT action type and only one (or very few) ItalWordNet senses. On many occasions it is possible to map a type onto a sense only excluding WordNet senses clearly referring to marked uses (metaphorical, idiomatic, etc.). With general verbs some difficulties emerge: sometimes an action type perfectly matches a WordNet sense or synset, but sometimes synsets are more generic than action types (and a best match may be found with hyponyms, if present). So the relations linking IMAGACT action types and ItalWordNet senses are the following: semantic equivalence, when a type perfectly matches a sense (ItalWordNet gloss perfectly describes the content of the video); otherwise, imperfect match, when the relation is one of subsumption (one type subsuming two or more senses, or two or more types being subsumed by one sense). We cannot exclude, a priori, the null relation (when a type cannot be related to any sense), but far we have not run into this.

Part of our future work will be to complete the mapping and to implement in IMAGACT, for each action type, an ItalWordNet direct link. We will also apply the same metodology to map English action types onto WordNet.

6   Conclusions and future work

The key innovation of IMAGACT is to provide a methodology which exploits the language independent capacity to appreciate similarities among scenes, distinguishing the identification of action types from their definition. By focusing its attention on action verbs, IMAGACT provides an interesting modality of presentation for their basic meaning distinctions; the navigation and search strategies are particularly promising for access to verbal meaning.

After its first delivery the IMAGACT infrastructure will grow freely as a function of its competence-based implementation in an open set of languages. The Interlinguistic Action Ontology DB will be available through the Internet as a web resource. The annotation infrastructure will be open source. We foresee that the infrastructure will have to cope with three main scenarios. The user may ask for:

a) the set of verbs of a target language that can be applied to a given action (language independent scenario);

b) the differential between the actions referred to by one verb in his own language and the actions referred to by a target verb in another language (distinguish the lexical properties of the target language in L2 acquisition);

c) the set of action types referred to by one or more action verbs in a given language (focusing on the lexical properties of action verbs).

The main NLP use foreseen for IMAGACT annotated data is word sense disambiguation. The resource will be tested in language acquisition and assisted translation scenarios; it will also be the starting point for the development of neuropsychological test batteries for the assessment of semantic knowledge[4]. Moreover the Ontology contains a large amount of information on actions potentially useful for ambient intelligence and for the modeling of artificial systems aimed at interacting in the natural environment on the basis of natural language instructions.

Acknowledgments

References

British National Corpus, version 3 (BNC XML Edition). 2007. Distributed by Oxford University Computing Services URL: http://www.natcorp.ox.ac.uk/

CLIPS Corpus. URL: http://www.clips.unina.it

C-ORALROM http://catalog.elra.info/product_info.php?products_id=757

Adamska-Sałaciak, A. 2008. Prepositions in Dictionaries for Foreign Learners: A Cognitive Linguistic Look. Bernal, E. and J. DeCesaris (Eds.). 2008: 1477-1485.

Cresswell M. F. 1978 Semantic Competence in F. Guenthner, M. Guenthner-Reutter, Meaning and translation. NY University Press: New York, 9-28

---

[4] There is a PhD thesis to be written on this topic.

De Mauro T., Mancini F., Vedovelli M., Voghera M. 1993. Lessico di frequenza dell'italiano parlato (LIP). Milano: ETASLIBRI.

Fellbaum, Ch. (ed.) 1998. WordNet: An Electronic Lexical Database. Cambridge: MIT Press.

Lew, Robert. 2010. 'New ways of indicating meaning in electronic dictionaries: hope or hype?' In: Zhang, Yihua (ed.), Learner's Lexicography and Second Language Teaching Shanghai: Shanghai Foreign Language Education Press. 387-404.

Rosch, E. 1978. Principles of Categorization. In E. Rosch & B.B. Lloyd (eds), Cognition and Categorization. Hillsdale: Lawrence Erlbaum, 27–48.

Wittgenstein, L. 1953. Philosophical Investigations. Oxford: Blackwell.