

An Exploratory Comparison of Thematic Roles in VerbNet and LIRICS

Claire Bonial* Susan Windisch Brown* William Corvey* Volha Petukhova**
Martha Palmer* Harry Bunt**

*Department of Linguistics, University of Colorado at Boulder

**Department of Computer Science, Tilburg University

{Claire.Bonial, Susan.Brown, William.Corvey, Martha.Palmer}@colorado.edu
{V.Petukhova, Harry.Bunt}@uvt.nl

Abstract

This research compares several of the thematic roles of VerbNet (VN) to those of the Linguistic InfRAstructure for Interoperable ResourCes and Systems (LIRICS). The purpose of this comparison is to develop a standard set of thematic roles that would be suited to a variety of natural language processing applications. Differences between the two resources are discussed, and in some cases, VN plans to adopt a corresponding LIRICS thematic role. In other cases, the motivations behind maintaining different thematic roles from those found in LIRICS are addressed.

1 Introduction

1.1 Motivation

The ideal set of thematic roles should be able to concisely label the arguments of any relation; however, what this set of roles should be has long been a subject of dispute in the linguistic community. In our current endeavor to create a possible standard set of thematic roles for the International Standards Organization (ISO), we have undertaken a systematic comparison of two semantic resources: LIRICS¹ and VerbNet (VN) (Schuler, 2002). We take a bottom-up approach in comparing roles across resources; the initial findings of this comparison follow, and demonstrate some of the difficulties in determining the ideal mapping between the thematic roles of LIRICS and VN.

1.2 Overview of LIRICS

The EU-funded project LIRICS was set up as a spin-off of ISO TC 37/SC4, with the aim of exploring the possibility of establishing sets of annotation concepts, defined in accordance with ISO

standard 12620 as so-called data categories, for syntactic, morphosyntactic, and semantic annotation and lexical markup. In the part of the project concerned with semantic annotation, several approaches and existing annotation schemes for semantic role labeling were analyzed and compared with respect to (1) description model; (2) granularity; (3) definition of semantic roles; and (4) consistency and reliability of annotation (Petukhova and Bunt, 2008). Based on this study, it was concluded that semantic roles should be defined:

- as neither syntactic nor lexical structures but as semantic categories;
- by virtue of distinctive semantic properties;
- that are not restricted to only a few specific verb (noun, adjective) classes;
- as relational notions that link participants to an event, describing the way the participant is involved in an event (e.g. does he act intentionally; is he/it affected, changed, manipulated by other participants; does it come into existence through the event), rather than by internal properties).

A set of 29 semantic roles² was defined by listing for each a characteristic set of entailments. These entailments were converted into a set of orthogonal properties, e.g. [+/- intentionality], [+/- independent existence], etc. (see also (Dowty, 1991) and (Sowa, 2000)). For example, the *Theme* role is defined as a participant in an event or state who (i) is essential to the event taking place but does not have control over the way the event occurs; (ii) is not structurally changed by the event; in a state, is in a fixed position or condition throughout the state; (iii) is causally involved or

¹Linguistic InfRAstructure for Interoperable ResourCes and Systems <http://LIRICS.loria.fr>

²This set includes 11 roles which are central to any event, e.g. *Agent, Theme, Patient*; 10 adjunct roles, e.g. *Time, Location, Manner*; and 8 sub-roles for *Time* and *Location*, e.g. *Duration, Frequency, Path*. For definitions and examples of see http://let.uvt.nl/general/people/bunt/docs/LIRICS_semrole.htm

affected by other participants; (iv) in a state is essential to the state being in effect; but it is not as central to the state as a participant in the Pivot role.

Different levels of granularity are distinguished, where a low-level semantic role inherits all the properties of a high-level role and has an additional feature, which reflects additional or more specific entailments.

The LIRICS set of semantic roles was evaluated with respect to redundancy, completeness and usability for reliable, consistent annotation using a multilingual test suite including English, Dutch, Italian and Spanish (see (Petukhova and Bunt, 2008) and (Bunt et al., 2007)).

1.3 Overview of VN

The purpose of VN is to classify English verbs based on semantic and syntactic regularities; it has been used for numerous NLP tasks, most notably, semantic role labeling ((Schuler , 2002) and (Shi and Mihalcea, 2005)). In each verb class, the thematic roles are used to link syntactic alternations to semantic predicates, which can serve as a foundation for further inferencing. For this reason, VN relies to an extent on syntactic features. Because VN is organized into verb classes, it is desirable to have an explicit hierarchy of roles such that users can understand the specificity of a role for a given class, as well as the superordinate category of that role, which would apply to classes of verbs that take diverse arguments. For example, the VN role *Topic* is a type of *Theme* that is restricted to arguments that express the transfer of information. The specificity of this role helps distinguish certain classes of verbs from others, and its compatibility with a particular verb helps determine whether that verb belongs in a certain class. However, thematic roles alone do not determine class membership; rather, a verb's thematic roles are considered along with the verb's semantics and syntactic patterning in assigning the verb class. The use of roles that are specific to certain classes of verbs is informative for VN users: roles that are unique to a particular class of verbs are maximally specific in their characterization yet amenable to hierarchical arrangement, which allows users to assign roles at various levels of granularity.

2 Thematic Roles in Comparison

LIRICS and VN thematic roles largely overlap; however, the divergent goals and structures of the

resources occasionally yield different roles. Attempting to find the ideal mapping between the roles of these two resources will be a first step in establishing an optimal set of standard thematic roles: roles that can generalize across the greatest number of syntactic and pragmatic contexts, while bringing the most appropriate level of specificity when naming an event participant. The following sections detail some of the challenges discovered in our initial comparisons of VN and LIRICS semantic roles.

2.1 VN *Actor1*, *Actor2* vs. LIRICS *Agent*, *Partner*

Verbs such as *chat*, *cooperate*, and *speak* correspond to events that usually involve two volitional participants, as in: 'Susan chatted/cooperated with Rachel.' Currently, VN uses the labels *Actor 1* and *Actor 2* to refer to each of these participants. In typical usage, *Actor 1* is the subject of the verb and *Actor 2* occurs in the oblique (e.g. 'with Rachel'). In theory, these labels capture the notion of two volitional actors involved in a single event, where one seems to be a true agent with pragmatic focus (*Actor 1*), while the other participant (*Actor 2*) fulfills the same agentive qualities (animate, volitional) without pragmatic focus.

While LIRICS does not have an exact mapping to *Actor 1* and *Actor 2*, it does have the complementary roles of *Agent* and *Partner*. In the LIRICS framework, an *Agent* is defined as a 'participant in an event who initiates and carries out the event intentionally or consciously, and who exists independently of the event,' while a *Partner* is defined as a 'participant in an event who is intentionally or consciously involved in carrying out the event, but who is not the principal agent of the event, and who exists independently of the event.' Upon examining this distinction between *Agent* and *Partner*, we decided that we preferred the LIRICS terms for the following reasons: 1) the labels *Agent* and *Partner* more clearly indicate that there are differing levels of agency between the two roles 2) using the term *Actor 1* fails to illustrate that the argument is essentially an agent.

2.2 VN *Theme 1*, *Theme 2* vs. LIRICS *Theme*, *Pivot*

Unfortunately, an adoption of *Agent* and *Partner* produces a potentially confusing incongruency among VN roles: parallel to *Actor 1* and *Actor 2*, VN has the roles *Theme 1*, *Theme 2*,

Patient 1 and *Patient 2*. *Theme 1* and *Theme 2*, for example, are used for verbs such as *border*, *coincide*, and *have*, which denote events that may involve two themes: ‘Italy-*Theme 1* borders France-*Theme 2*.’ The relationship between the two themes is analogous to the relationship between *Agent* and *Partner*: there is a pragmatically focused theme (*Theme 1*) and a secondary theme (*Theme 2*). In order to accommodate this parallelism, we decided to maintain the concept behind LIRICS *Agent* and *Partner*, but adjust the labels to *Agent* and *Co-Agent*, *Theme* and *Co-Theme* and *Patient* and *Co-Patient*. Perhaps more importantly, there is the rare possibility that a sentence could involve both a *Partner* to an *Agent* and a *Partner* to a *Theme* or *Patient* leading to two ambiguous *Partner* arguments; the ‘Co-’ terminology allows them to be easily distinguished.

In certain cases, we found that the role *Theme 1* could be better expressed using another LIRICS role: *Pivot*, a ‘participant in a state that is characterized as being in a certain position or condition throughout the state, and that has a major or central role or effect in that state.’ For verbs in the Own and Require classes, using *Theme 1* to refer to the possessor or requirer seemed to obscure an important distinction between this type of participant and other *Theme 1* arguments, wherein the *Theme 1* is primarily being located (e.g. ‘Italy’ in ‘Italy borders France’). For verbs in the Own and Require classes, *Theme 1* is not located; instead it is involved in a state of ownership or need. Thus, for the Own and Require classes, we did not adopt *Theme* and *Co-Theme* to replace *Theme 1* and *Theme 2*; rather, we chose to adopt the label *Pivot* for participants in a state of ownership or need, and to use *Theme* to refer only to the owned or needed participants. We expect to utilize *Pivot* in similar circumstances throughout the resource.

2.3 VN *Topic* vs. LIRICS *Theme*

In VN, *Theme* is used with a wide variety of verbs to label a participant that is being literally or metaphorically located, positioned, or moved; this participant may be concrete or abstract. *Topic*, on the other hand, is restricted to participants involved in the transfer of information: arguments of verbs such as *advise*, *promise*, and *tell*. For example: ‘John-*Agent* informed me-*Recipient* of the situation-*Topic*.’ *Topic* inherits all of the features of *Theme*, but is constrained by additional features

such as +*information content* and +*abstract*. LIRICS does not use *Topic*, instead *Theme* would be used for these arguments. This discrepancy illustrates the differing aims of the two resources: VN uses finer-grained roles where this can help to distinguish classes, a practice LIRICS specifically avoids.

2.4 VN *Stimulus*, *Experiencer* vs. LIRICS *Cause*, *Pivot*

Verbs such as *see*, *amuse* and *empathize* involve one participant that is perceiving another in a cognitive or sensory manner, but the event does not necessarily involve contact or volition on the part of either party. The participant that triggers the event does not do so purposefully; the fact of its existence, perceived by another participant, yields the physical or mental reaction in that participant. In the LIRICS framework, the trigger would be a *Cause*, defined as a ‘participant in an event (that may be animate or inanimate) that initiates the event, but that does not act with any intentionality or consciousness; it exists independently of the event,’ while the participant reacting would be *Pivot*. In VN, the trigger is inconsistently labeled either *Cause* or *Stimulus*, while the participant reacting is the *Experiencer*. After an examination of all of the verb classes using these roles in combination (27 classes), we defined *Stimulus* as a participant that unintentionally arouses a mental or emotional response in a sentient being. In turn, we found that the *Experiencer* was consistently a participant undergoing a particular mental or emotional state precipitated by the mere perception of another participant. Thus, we found that *Stimulus* and *Experiencer* emerged as a natural pairing in verb classes involving a cognitive or emotional event. *Stimulus* is thus a more constrained type of *Cause*, where the causation is mediated by cognitive experience: ‘The storm-*Stimulus* frightened the children-*Experiencer*,’ vs. ‘The storm-*Cause* destroyed the ship-*Patient*.’ Therefore, although we find the LIRICS roles of *Pivot* and *Cause* to be very useful and have their place within VN, we also believe that the greater specificity of the *Experiencer* and *Stimulus* roles, which helps to distinguish verb classes, should be maintained.

2.5 Discussion: Remaining open issues

As demonstrated in the comparisons presented, decisions concerning one thematic role often impact other thematic roles and thematic role pat-

terns across the resource. For this reason, it is important to keep multiple thematic roles in mind when analyzing the impact of proposed changes. Further, this paper has explored possible subset relations among thematic roles, raising questions about the nature and depth of hierarchical relationships among roles. For instance, the LIRICS role *Goal* corresponds to VN *Recipient*, and the LIRICS *Final Location* corresponds to VN *Destination* (Petukhova and Bunt 2008). However, in many semantic frameworks, *Goal* represents an end location of an action that would subsume both *Recipient* and *Destination*. In examining possible benefits of incorporating *Goal* into a hierarchy of VN roles, we confront questions about congruency of scope (in this case, *Goal* versus *Final Location*) between semantic roles at a given level within a hierarchy.

Analysis of several VN roles is still underway. As we have begun to show above, the roles *Source*, *Location*, *Destination*, and *Recipient* are closely related to each other, and an initial look into their use in VN suggested that the definitions should be clarified and additional roles considered. For example, the *Destination* role was initially used for goals that were physical locations but had been extended as new classes were added to include non-locative goals. In addition, its use seemed to overlap in some cases with the role *Location*. Our comparison with the LIRICS roles *Location*, *Initial Location*, *Final Location*, *Source*, *Goal*, and *Recipient* is contributing to our construction of explicit definitions for these VN roles and our consideration of new roles. The LIRICS features concerning the temporality and physical locality of these roles is helping direct our analysis. Additionally, the status of the role *Proposition* is in question, as its distinction from *Topic* may be purely syntactically motivated (i.e. *Propositions* only occur as clausal arguments). A summary of the role comparisons completed in the present study appears in Table 1.

2.6 Conclusion and Future Work

In this comparison process we are re-evaluating VN roles, allowing us an opportunity to create a clear definition for each role and to make changes ensuring that each role is used consistently throughout VerbNet; these definitions and changes are forthcoming. Our ongoing comparison of VN and LIRICS has demonstrated that

Current VN role	LIRICS role	Proposed VN role
Actor 1	Agent	Agent
Actor 2	Partner	Co-agent
Patient 1	Patient	Patient
Patient 2	Partner	Co-patient
Theme 1	Pivot	Theme
Theme 2	Theme	Co-theme
Theme 1	Pivot	Pivot
Theme 2	Theme	Theme

Table 1: Summary of VerbNet role changes based on comparison to LIRICS (note that the exact *Theme 1* and *Theme 2* changes will depend upon the verb class under consideration).

resources differing in aim and structure can still overlap a great deal in their definitions of core thematic roles (e.g., *Agent*, *Patient*, *Instrument*, and to a large extent, *Theme*). Their differing goals can also result in some variation in their final sets of thematic roles and the boundaries of those roles. Although this highlights the difficulties that are involved in creating an ISO standard set of thematic roles, the process of comparison has also made explicit the motivations behind certain differences. In some cases, the comparison led to a revision of VN roles (e.g., adopting *Pivot* in place of *Theme* in certain situations, and changing *Actor 1* and *Actor 2* to *Agent* and *Co-agent*), whereas in others the comparison helped develop more rigorous role definitions (e.g., for *Experiencer* and *Stimulus*). With a clearer understanding of the resources’ motivations for the roles, we are better able to devise a set of thematic roles that are suited to the widest range of purposes.

Future work will perform similar comparisons over additional resources, notably FrameNet (Fillmore and Baker, 2010), as LIRICS already defines links to roles in this resource. The final goal of all comparisons will be the development of a set of thematic roles that is suited not just to the idiosyncratic purposes of one resource, but rather to a wide variety of natural language processing purposes. As we have shown in these initial comparisons, one of the most difficult issues in developing a standard resource compatible with different purposes is the issue of granularity, or the extent to which thematic roles are illustrative of different classes of verbs as opposed to generalizable across all verbs. To overcome this difficulty, we are adjusting VN such that the resulting thematic role

set will be hierarchical. A hierarchical structure with clear mappings between higher and lower-order classes will allow users to select the level of granularity that is best suited to the application they are developing, including only those thematic roles that generalize across all classes of verbs (e.g. *Theme*), or including more specific roles that are characterized by additional features, and therefore only appear in certain classes of verbs (e.g. *Topic*).

Research in automatic semantic role labeling, for example, has demonstrated the importance of the level of granularity of semantic roles. Yi, Loper and Palmer (2007) and Loper et al. (2007) both demonstrate that because VN labels are more generalizable across verbs than PropBank (Palmer et al., 2005) labels, they are easier for semantic role labeling systems to learn. However, Merlo and Van Der Plas (2009) found that the differing levels of granularity of PropBank and VN were both useful, and therefore suggest complementary use of both resources. Our hope is that the final set of thematic roles we decide upon, informed by our comparisons to other resources such as LIRICS and FrameNet, will encompass the benefits of hierarchical granularity, thereby meeting the unique needs of varying natural language processing applications.

Acknowledgments

We gratefully acknowledge the support of the National Science Foundation Grant NSF-0415923, Word Sense Disambiguation, and a DARPA supplement to a grant from the Defense Advanced Research Projects Agency (DARPA/IPTO) under the GALE program, DARPA/CMO Contract No. HR0011-06-C-0022, a subcontract from BBN, Inc.

References

Bunt, H., Petukhova, V., and Schiffrin, A. 2007. *LIRICS Deliverable D4.4. Multilingual test suites for semantically annotated data*. <http://lirics.loria.fr>.

Dowty, D. 1991. Thematic Proto-Roles and Argument Selection. *Language*, 67:547-619.

Fillmore, C.J. and Baker, C.F. 2010. A Frame Approach to Semantic Analysis. In Heine, B. and Narrog, H. (eds.) *Oxford Handbook of Linguistic Analysis*: Oxford University Press.

Kipper, K. 2002. VerbNet: A Class-Based Verb Lexicon. <http://verbs.colorado.edu/mpalmer/projects/verbnnet.html>

Loper, E., Yi, S., and Palmer, M. 2007. Combining Lexical Resources: Mapping Between PropBank and VerbNet. Proceedings of the Seventh International Workshop on Computational Semantics (IWCS-7), Tilburg.

Merlo, P., and Van Der Plas, L. 2009. Abstraction and Generalization in Semantic Role Labels: PropBank, VerbNet or both? Proceedings of the 47th Annual Meeting of the ACL and 4th IJCNLP of the AFNLP, Suntec, pp. 288-296.

Palmer, M., Gildea, D., and Kingsbury, P. 2005. The Proposition Bank: An annotated Corpus of Semantic Roles. *Computational Linguistics*, 31:1, pp. 71-105.

Petukhova, V., Schiffrin, A., and Bunt, H. 2007. Defining Semantic Roles. Proceedings of the Seventh International Workshop on Computational Semantics (IWCS-7), Tilburg, pp. 362-365.

Petukhova, V., and Bunt, H. 2008. *LIRICS semantic role annotation: Design and evaluation of a set of data categories*. Proceedings of the sixth international conference on language resources and evaluation (LREC 2008), Paris: ELRA.

Schuler, K.K. 2005. *VerbNet: A Broad-Coverage, Comprehensive Verb Lexicon* University of Pennsylvania.

Shi, L. and Mihalcea, R. 2005. *Putting pieces together: Combining FrameNet, VerbNet and WordNet for robust semantic parsing*. *Computational Linguistics and Intelligent Text Processing*, 100-111.

Sowa, J.F. 2000. *Knowledge representation: logical, philosophical and computational foundations*. Pacific Grove: Brooks/Cole.

Swier, R.S. and Stevenson, S. 2004. Unsupervised semantic role labelling *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, 95-102.

Yi, S., Loper, E., and Palmer, M. 2007. Can semantic roles generalize across genres? Proceedings of the HLT/NAACL-2007, Rochester, pp. 548-555.