
**Language resource management —
Semantic annotation framework
(SemAF) —**

**Part 4:
Semantic roles (SemAF-SR)**

*Gestion de ressources linguistiques — Cadre d'annotation sémantique
(SemAF) —*

Partie 4: Rôles sémantiques (SemAF-SR)



STANDARDSISO.COM : Click to view the full PDF of ISO 24617-4:2014



COPYRIGHT PROTECTED DOCUMENT

© ISO 2014

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized otherwise in any form or by any means, electronic or mechanical, including photocopying, or posting on the internet or an intranet, without prior written permission. Permission can be requested from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office
Case postale 56 • CH-1211 Geneva 20
Tel. + 41 22 749 01 11
Fax + 41 22 749 09 47
E-mail copyright@iso.org
Web www.iso.org

Published in Switzerland

Contents

Page

Foreword	iv
Introduction	v
1 Scope	1
2 Terms and definitions	1
2.1 Formal semantic units	1
2.2 Conceptual semantic units and relations	2
2.3 General linguistic units	3
3 Abbreviated terms	4
4 Purpose and justification	4
5 The nature of semantic roles	4
5.1 General	4
5.2 Typical examples of semantic roles	6
6 Metamodel	7
6.1 Key concepts	7
6.2 Introduction to a metamodel for semantic role annotation	11
6.3 Criteria for distinguishing semantic roles	12
6.4 Defining eventuality frames	13
7 Interactions	14
7.1 Semantic roles and semantic types	14
7.2 Complexities	14
7.3 The role of context in semantic role assignment	14
7.4 Fuzzy boundaries between roles	15
7.5 Multiple classification	15
7.6 Inheritance relations between semantic roles	15
8 Guidelines for developing new semantic role frameworks for languages and/or domains	16
8.1 General	16
8.2 Mapping from VerbNet to LIRICS	17
Annex A (informative) Specification of ISO semantic roles	19
Annex B (informative) Review of existing frameworks	28
Annex C (informative) Specification of the annotation language	37
Bibliography	41

Foreword

ISO (the International Organization for Standardization) is a worldwide federation of national standards bodies (ISO member bodies). The work of preparing International Standards is normally carried out through ISO technical committees. Each member body interested in a subject for which a technical committee has been established has the right to be represented on that committee. International organizations, governmental and non-governmental, in liaison with ISO, also take part in the work. ISO collaborates closely with the International Electrotechnical Commission (IEC) on all matters of electrotechnical standardization.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular the different approval criteria needed for the different types of ISO documents should be noted. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (see www.iso.org/directives).

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO shall not be held responsible for identifying any or all such patent rights. Details of any patent rights identified during the development of the document will be in the Introduction and/or on the ISO list of patent declarations received (see www.iso.org/patents).

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation on the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the WTO principles in the Technical Barriers to Trade (TBT) see the following URL: Foreword - Supplementary information

The committee responsible for this document is ISO/TC 37, *Terminology and other language and content resources*, Subcommittee SC 4, *Language resource management*.

ISO 24617 consists of the following parts, under the general title *Language resource management — Semantic annotation framework (SemAF)*:

- Part 1: Time and events (*SemAF-Time, ISO-TimeML*)
- Part 2: Dialogue acts
- Part 4: Semantic roles (*SemAF-SR*)
- Part 5: Discourse structure (*SemAF-DS*)
- Part 7: Spatial information (*ISO-Space*)

The following parts are under preparation:

- Part 8: Semantic relations in discourse (*SemAF-DRel*)

Principles of semantic annotation (SemAF-Basics) will form the subject of future Part 6.

Introduction

This part of ISO 24617 aims to specify criteria for defining semantic roles (SRs), and is the outcome of an agreement that the various semantic role frameworks being used to support data annotation (e.g. FrameNet, VerbNet, PropBank, EngVallex, and LIRICS, to name only a few examples for English) have strong underlying compatibilities. The goal is to provide both an explanation of these compatibilities and a loose mapping between definitions of individual semantic roles, as listed in the different frameworks, that will benefit the community as a whole.

The current specification has been developed under the aegis of the ISO Semantic Annotation Framework (SemAF), where it is known as SemAF-SR.

The main parts of ISO 24617-4 consist of the following:

- Scope;
- Normative references;
- Terms and definitions;
- motivation and requirements;
- basic concepts and metamodel specifications;
- examples of mapping existing frameworks to the metamodel.

This part of ISO 24617 contains three informative annexes. In [Annex A](#), the ISO semantic roles are specified. In [Annex B](#), information is provided both on past and current activities in semantic role annotation and on tools and frame files. [Annex C](#) contains the abstract and concrete syntax for the metamodel.

STANDARDSISO.COM : Click to view the full PDF of ISO 24617-4:2014

Language resource management — Semantic annotation framework (SemAF) —

Part 4: Semantic roles (SemAF-SR)

1 Scope

The aim of this part of ISO 24617 is to propose a consensual annotation scheme for semantic roles; that is to say, a scheme that indicates the role that a participant plays in an event or state, as described mostly by a verb, and typically providing answers to questions such as “‘who’ did ‘what’ to ‘whom’”, and ‘when’, ‘where’, ‘why’, and ‘how’. This includes not only the semantic relations between a verb and its arguments but also those relations that are relevant for other predicative elements such as nominalizations, nouns, adjectives, and predicate modifiers; the predicating role of adverbs and the use of coercion fall outside the scope of this part of ISO 24617.

NOTE In linguistics, **coercion** occurs when the grammatical context causes the language-user to reinterpret all or parts of the semantic and/or formal features of a lexeme that appear in that context.^[60]

2 Terms and definitions

For the purposes of this document, the following terms and definitions apply.

2.1 Formal semantic units

2.1.1

argument

formal semantic unit that is an essential element of a *predicate argument structure* (2.1.3) and can have variable instantiations depending on the utterance

Note 1 to entry: An argument corresponds to a *participant* (2.2.5) of an *eventuality* (2.2.2) described by the *predicate argument structure* (2.1.3).

Note 2 to entry: Arguments typically satisfy certain argument positions and can be described as being syntactico-semantic notions, whereas *participants* (2.2.5) are semantico-conceptual. The standard view is that subsets of the *participants* associated with an *eventuality* (2.2.2) are selected as arguments by the verb (or nominal or adjective) expressing the *eventuality* (2.2.2). Other *participants* (2.2.5) are either incorporated or realized as *eventuality modifiers* (2.2.4).

Note 3 to entry: Natural language predicates typically have one, two, or three arguments, although they can have more.

2.1.2

predicate

formal semantic unit that represents a semantic relation between one or more *arguments* (2.1.1) in a *predicate argument structure* (2.1.3)

Note 1 to entry: Predicates are indicated by predicative linguistic elements such as verbs, nouns, and adjectives.

2.1.3

predicate argument structure

formal representation of the core semantic content of an utterance, consisting of a *predicate* (2.1.2) constant, and its *arguments* (2.1.1)

Note 1 to entry: In classical logic-based semantics, this corresponds to predicate argument structures in first-order predicate logic.

Note 2 to entry: One of the *arguments* (2.1.1) can be a variable uniquely identifying the instance of the predicate argument structure to allow references to it in other predicate argument structures.

Note 3 to entry: The representation of event semantics is subject to many variations; some of them, such as in Reference [41], can have separate *predicates* (2.1.2) for each *semantic role* (2.2.6) relation. In this case, the predicate argument structure of an utterance is the sum of the individual predicate *semantic role* (2.2.6) assertions representing the semantic content of the utterance.

2.2 Conceptual semantic units and relations

2.2.1

entity

conceptual semantic unit that typically functions as a *participant* (2.2.5)

Note 1 to entry: An entity is an individual such as a person, organization, physical object, or logical entity, as well as, on occasion, a number, quantity, dimension, or a reification of an eventuality, a property, or a quality, e.g. emotion (anger, love), the value of a colour, etc.

2.2.2

eventuality

event, state, process, or action which can have *participants* (2.2.5) and which is being referred to by a verbal, adjectival, or nominal description in an utterance

Note 1 to entry: The formal representation of an eventuality is a *predicate argument structure* (2.1.3).

Note 2 to entry: See ISO 24617-1. An eventuality can also be described as 'something that can be said to obtain or hold true, to happen or to occur', as in ISO 24617-1. As such, they can be actual, hypothetical, or generic, covering situations such as "You should go home," or "He might be John's brother."

2.2.3

eventuality frame

generalized abstract specification of the *word sense* (2.3.6) associated with an *eventuality* (2.2.2) in an utterance

Note 1 to entry: The frame consists of the specification of (a) a *predicate* (2.1.2) that can participate in a class hierarchy if such a hierarchy is specified, and (b) the *arguments* (2.1.1) that this *predicate* (2.1.2) expects along with their *semantic roles* (2.2.6).

2.2.4

eventuality modifier

particular type of *participant* (2.2.5) that completes the description of an *eventuality* (2.2.2) but is optional and not essential

Note 1 to entry: Eventuality modifiers are distinct from other types of participants in that they are used in supplying information that is typically more peripheral and more general, for example, situating the eventuality in time or space.

Note 2 to entry: In FrameNet, these would be peripheral frame elements and in PropBank, ArgM's.

Note 3 to entry: Eventuality modifiers typically correspond to syntactic adjuncts.

2.2.5**participant**

conceptual semantic unit referred to by one or more lexical items in an utterance, which is or can be involved in an *eventuality* (2.2.2)

Note 1 to entry: Both *entities* (2.2.1) and *eventualities* (2.2.2) can function as participants.

2.2.6**semantic role**

mode of involvement of a *participant* (2.2.5) in an *eventuality* (2.2.2)

Note 1 to entry: Semantic roles for specific eventualities are often associated with prototypical semantic relations, e.g. if *John* causes a *breaking* event, he is the *Agent*; if he uses a *hammer*, it is the *Instrument*; and someone who *receives* something is a *Recipient* (see [Clause 5](#) for descriptions).

2.3 General linguistic units**2.3.1****lemma****lemmatized form**

conventional form chosen to represent a *lexeme* (2.3.2)

Note 1 to entry: See ISO 24611.

2.3.2**lexeme**

fundamental unit, generally associated to a set of word forms sharing a common meaning

Note 1 to entry: See ISO 24611.

2.3.3**lexical entry**

container for managing a set of word forms and possibly one or several meanings [*word senses* (2.3.6)] to describe a *lexeme* (2.3.2)

Note 1 to entry: See ISO 24611.

2.3.4**lexicon**

resource comprising a collection of *lexical entries* (2.3.3) for a language

Note 1 to entry: See ISO 24611.

2.3.5**utterance**

stretch of speech about which no assumptions have been made in terms of linguistic theory

Note 1 to entry: See Reference [12].

2.3.6**word sense**

meaning associated with a *lexeme* (2.3.2) in a context

Note 1 to entry: The 'river bank' sense of *bank* and the 'financial institution' sense of *bank* are considered to be two different word senses, or *lexical units*, with the same word form, or *lexeme* (2.3.2). *I called him on the radio* and *Call me a taxi* are associated to different word senses of the *lexeme* (2.3.2) *call*. Unrelated senses, as in *bank*, are called *homonyms*. Senses of the same word form or *lexeme* which are clearly related (and can be difficult to distinguish) are called *polysemes*, e.g. *Coins with an image of the king, preoccupied with body image, evokes a strong mental image*.

3 Abbreviated terms

EngVallex	English Valency Lexicon
LIRICS	Linguistic Infrastructure for Interoperable Resources and Systems
PropBank	Proposition Bank
SR	semantic roles
SRL	semantic role labelling
WSD	word sense disambiguation

4 Purpose and justification

Semantic roles are arousing increasing interest in the information-processing community because they make explicit the key conceptual relations of participation between a verb and its arguments; that is to say, they specify 'who' did 'what' to 'whom', and 'when', 'where', 'why', and 'how'. For English alone, there are already several different semantic role frameworks, including FrameNet, VerbNet, LIRICS, EngVallex, and PropBank. Although these have been developed independently, there are strong underlying compatibilities between them, and they share a central definition of what a semantic role is, and what its span is, within an individual sentence. In addition to defining key concepts, this part of ISO 24617 aims to clarify and specify these underlying compatibilities and provide, where possible, a mapping between similar semantic roles across different frameworks. This mapping is intended to serve as an illustration of how different semantic role definitions can be linked to each other across frameworks, and presupposes a specification of clearly defined criteria for distinguishing semantic roles.

The specification will be used in two different situations:

- in annotations where the semantic roles are recorded in annotated corpora;
- as a dynamic structure produced by automatic systems, a process typically known as semantic role labelling (SRL).

The objectives of this specification are to provide

- a reference set of data categories that define a structured collection of semantic roles with an explicit semantics,
- a pivot representation based on a framework for defining semantic roles that can facilitate mapping between different formalisms (alternative semantic role representations/syntactic theories/eventually different languages) and, in the future, between different languages, and
- guidelines for creating new resources that will be immediately interoperable with pre-existing resources.

5 The nature of semantic roles

5.1 General

For computers to make effective use of information encoded in text, they must be able to detect the eventualities that are being described and the eventuality participants. The processing of a sentence like *John threw a ball to Mary in the park* should result in the identification of a throwing event involving *John* as the Agent of the event, *Mary* as the Recipient, and the *ball* as the item being *thrown*; the location of the throwing event, or where it took place, is *the park*. This description of the event specifies the conceptual relations of participation that the referents of the noun phrases play with respect to the event. The semantic notions being specified are the *roles of the participants in an eventuality (i.e. semantic roles)*.

This part of ISO 24617 establishes LIRICS (see [Annex A](#)) as a reference set of semantic roles with precise definitions. Researchers are free to define their own sets of semantic roles, but explicit information on how they can be mapped to the reference set will make resources more interoperable. Many resources currently map to PropBank, VerbNet, or FrameNet. Since this part of ISO 24617 includes mappings of these resources to LIRICS, such mappings already qualify as meeting the requirement of interoperability.

Our *throw* example seems fairly straightforward, but complexities quickly arise. English, for instance, allows not only several different syntactic constituents to present the same semantic role, but also several different semantic roles to be presented by the same syntactic constituent. For decades, a central concern of linguists has been the elucidation of the process of mapping back and forth between the syntactic analysis of the sentence and the conceptual structure and relations in the event described. For example, in the following two sentences,

(1) *The flame melted the wax.*

(2) *The wax melted.*

a standard syntactic parser represents the *wax* as the verb's direct object in the first sentence and its subject in the second. There is nothing overt to indicate that it has the same conceptual relation in both cases despite the fact that it is expressed syntactically in a different way. We can capture this by annotating the *wax* as having the same semantic role (or conceptual relation) in both sentences. It would typically be labelled the Patient, the participant undergoing a change of state. Note that both sentences are in the active voice, not the passive voice. In *The wax was melted by the flame*, the passive provides syntactic evidence that the *wax* is playing the same role (Patient) that it plays in example (1). Since the particular pair of syntactic variations illustrated by *melt* does not occur with every transitive verb [see example (5)], it is not easily predictable. Other transitive verbs can also occur in an intransitive form while maintaining the same semantic role for the subject as the transitive, as in the following example, where *soprano* is the Agent of *sing* in both sentences (the *aria* is the Theme):

(3) *The soprano sang an aria.*

(4) *The soprano sang.*

The verb *slice* can also move the Patient (*the bread*) to subject position, as in

(5) *John sliced the bread easily. / This bread slices easily.*

although other transitive verbs, such as *eat*, cannot:

(6) *John ate the apple.*

(7) *John ate.*

(8) **The apple ate crunchily in the background.*

The last sentence is starred (*) to indicate its ungrammaticality.

Accurate interpretation of the semantic roles of the verb arguments (i.e. 'Who did what to whom?') is a crucial goal for natural language processing systems. Our ability to do this automatically has improved enormously in recent years and has been largely based on the availability of annotated corpora. In fact, there are corpora, such as FrameNet and PropBank, available with quite different semantic role annotations, and this prompts questions about the nature and number of semantic roles. This part of ISO 24617 attempts to provide definitions and examples clarifying their definition.

For semantic roles to maximize the benefit to the information processing community, it is desirable that the definitions of the semantic roles should, as far as possible, have the following properties:

- consistently recognizable;
- able to clarify sense distinctions;
- generalizable;

- machine learnable;
- able to provide an appropriate foundation for inferencing.

The purpose of the specifications for semantic role definitions in this part of ISO 24617 is to provide these attributes.

5.2 Typical examples of semantic roles

A list of the best-known roles and the properties usually associated with them is given below. They are taken from the EAGLES discussion on Standardizing Subcategorization (see Reference [61]). Comments in parentheses have been added for clarification purposes and include comparisons with LIRICS (Linguistic Infrastructure for Interoperable Resources and Systems), and occasionally with VerbNet.

NOTE These role definitions, and the ones for LIRICS, are quite general and can cover a wide range of participant types. The specific preposition associated with an individual semantic role quite often adds nuances of meaning in addition to that conveyed by the semantic role itself. For instance, in “*Eat the fish with caution*,” the phrase “with caution” would typically be labelled as MANNER. Another example of a MANNER role could be the phrase “in three bites” from “*Eat the fish in three bites*.” Clearly the interpretation of these two phrases will be quite different, as will be their impact on the representation of the *eating* eventuality. Semantic role labels cannot be expected to clarify all such types of subtle differences in meaning, and additional research on the definitions of prepositions and their interactions with semantic roles, such as Srikumar and Roth, 2013[48], is needed.

Agent

A participant that the meaning of the verb specifies as doing or causing something, possibly intentionally; for example, as the subject of *kill*, *eat*, *hit*, *smash*, *kick*, and *watch*. (LIRICS has a similar Agent, which acts intentionally or consciously.)

Patient

A participant that the verb characterizes as having something happen to it, and as being affected by what happens to it; for example, as the object of *kill*, *eat*, and *smash*, but not of *watch*, *hear*, and *love*. (If someone watches television, the television is not affected by the watching, so it would be a Theme rather than a Patient. LIRICS has a similar Patient.)

Experiencer

A participant that is characterized as aware of something; for example, as the subject of *love* or as the object of *annoy*. (LIRICS uses an Agent for these verbs and has no Experiencer; VerbNet has the same role, but only when the Experiencer is affected by the event.)

Theme

A participant that is characterized as changing its position or condition, or as being in a state or position; for example, as the object of *give* and *hand* or as the subjects of *walk* (in line with the policy of labelling the object in motion as a Theme) and of *die*. (According to the EAGLES definition, people who die of old age would not be considered to be Patients because an Agent has not acted upon them. EngVallex stays more syntactically oriented and marks such participants as Actors or Patients, depending on their syntactic position.)

Location

The thematic role associated with the Noun Phrase expressing the location in a sentence with a verb of location [perhaps in a Prepositional Phrase (PP)]; for example, the subjects of *keep*, *own*, *retain*, and *know*, and locative PPs. (For EAGLES, the location of the thing *kept* or *owned* is considered to be with the *keeper* or the *owner*, and so on. LIRICS has a Location role, but uses it in a more restrictive way, and uses Agent or Pivot for several of these verbs instead of Location, as does VerbNet.)

Source

Object from which motion proceeds; for example, as the subject of *sell* and *promise* or as the objects of *deprive*, *free*, and *cure*. Note that the motion can be abstract. (LIRICS uses Source for these abstract examples, but uses Initial-Location instead for physical motion verbs.)

Goal

Object to which motion proceeds (e.g. the Path-prepositional phrase for caused-motion verbs like *throw*, and for metaphorical motion events), the subject of *receive* and *buy* and the dative object of *tell* and *give*. (Adapted from Dowty, 1989.^[21] LIRICS has a for abstract objects and uses Final-Location for physical motion.)

Although these semantic roles were initially defined in relation to the arguments and syntactic adjuncts of verbs, many theories recognize semantic roles associated with nouns, particularly event nominals related to verbs. *The Romans, Jerusalem, and 70 C.E.* therefore play the same roles in both examples (9) and (10):

(9) *The Romans destroyed Jerusalem in 70 C.E.*

(10) *The destruction of Jerusalem by the Romans in 70 C.E. is famous.*

Nombank, which extends PropBank to nouns, handles this by using the same roleset for both the verb and the noun; FrameNet handles it by including *destroy* and *destruction* in the same semantic frame, and annotating each with the same frame elements.

6 Metamodel

6.1 Key concepts

The discussion of a metamodel begins by defining the following concepts:

- eventuality frames;
- arguments/adjuncts;
- granularity of word senses and semantic roles;
- semantic classes;
- entailment.

6.1.1 Eventuality frames

This part of ISO 24617 adopts the classic type-token distinction that is accepted by philosophical and knowledge representation communities, which is to say that a concept is distinct from the objects that are instances of it. With respect to eventualities, an individual utterance can correspond to an instance, or a token, of a particular eventuality frame (2.2.3). (Note that an utterance does not necessarily describe an eventuality, e.g. *Oh, no!* or *Ouch!*.) All utterances that describe eventualities involving objects being *thrown* can be considered to be tokens of the eventuality frame '*throwing of objects*'. Semantic role labelling frameworks require the creation of lexical entries for each eventuality frame to be annotated; these lexical entries are described at the type level and illustrated with several naturally occurring sentences that describe tokens. For instance, for each lexeme PropBank defines a Frame File with one or more framesets (also called rolesets) corresponding to individual eventuality frames (FrameNet associates it with a Frame, VerbNet with a class, and EngVallex with a valency entry). These are used to guide the annotation process.

6.1.2 Arguments and adjuncts

The argument/adjunct distinction is central to a definition of semantic roles because of the different relations that arguments and adjuncts (defined in [Clause 2](#) as eventuality modifiers) have with regard to the predicate. This distinction is not always easy to grasp, but in practice, nearly all successful semantic role annotation projects assume some variation of this distinction. Arguments are characterized by the following three properties. The first property is obligatoriness. A predicate is central to an eventuality description, but the arguments are the necessary elements without which the eventuality description is incomplete. This makes arguments obligatory, at least prototypically. Obligatoriness cannot be defined on purely syntactic grounds, since arguments that are essential to a predicate are routinely dropped in surface realizations; for example, in example (7), repeated as example (11), the thing that is *eaten by John* is not realized syntactically, but there is still a strong sense that it is an important element in the eventuality description, and a core role has to be assigned to this argument. By contrast, although the *eating* event must have occurred at a certain time and in a certain place, the time and place are not necessary for the eventuality description to be complete, and are considered adjuncts, or eventuality modifiers.

(11) *John ate.*

Being obligatory also seems to be linked to a second property of arguments, which is that the types of arguments associated with a predicate can be specific to that predicate since they are based on its semantics. In other words, different predicates tend to take different sets of arguments. A Recipient is a receiver of an item or a message, as in *He gave a book to Mary*. This role is natural in these sentences but would be very odd in the description of an *eating* eventuality. This contrasts with typical adjuncts like time and location, which can occur with a wide variety of quite different predicates. To put it another way, arguments tend to co-vary with the predicate while adjuncts do not. Statistically speaking, arguments are also more likely than adjuncts to co-occur with the predicate.

The third property of arguments is that they are often assumed to be unique, and multiple arguments are not expected to fill the same role and have the same semantic relationship with the predicate. A notable exception to this can occur with plurals, which for many verbs can also be split into separate argument slots. For example, *John and Mary met*, can also be readily phrased as *John met Mary*, and with such subtle differences, if any, in focus and agentivity that they are sometimes labelled Agent and Co-Agent. A significant number of 'reciprocal' predicates such as *fight*, *match*, *marry*, *play* (in the sense of playing a two-person game like tennis), and *tie* (to score equal points) share this property. Reciprocity is not confined to Agents and could also characterize Themes (*match*), and perhaps other roles as well.

Different semantic role annotation schemes, which are described in [Annex B](#), often choose to formalize all, or a subset, of these properties, as discussed in that annex. In theoretical linguistics, there is much controversy over where the boundary should be drawn between arguments and adjuncts. For a fairly radical view setting out a fluid characterization of the boundary between arguments and adjuncts, see Reference [15]. Generally speaking, arguments are considered to influence the interpretation of the verb while adjuncts are not.

6.1.3 Granularity of word senses and semantic roles

The semantic ambiguity of words is pervasive; many words, especially the most frequently used words, have multiple meanings. For example, one can *draw a gun*, *draw water from a well*, and *draw a diagram*. Depending on the dictionary, an entry for a word like *draw* or *run* can list dozens of separate definitions, but unless they encounter puns or other forms of word-play, people rarely have difficulty interpreting the meaning of these words. Although lexical ambiguity can present very few problems to people engaged in normal interpretation of text or speech, the same cannot be said for computers. The correct selection of the appropriate meaning of a word in context has proved to be very difficult for natural language processing (NLP) systems. Yet an accurate means of performing word sense disambiguation (WSD) would improve many NLP applications such as information extraction, information retrieval, and machine translation, and also any task that requires more complex knowledge representation and reasoning.[8] [9] [24] [47] [50] Encouraging progress has been made recently with automatic word sense disambiguation systems that approach human levels of accuracy.[68]

This is equally true for semantic role labelling. Different senses of the same word can refer to different event types and therefore require different semantic roles. For instance, the possible association of an Agent and a Co-Agent mentioned above only applies to a particular sense of *play*, as in *John and Bill played tennis* or *John played tennis with Bill*. The acting sense, as in *Olivier played Hamlet* or in the musical performance sense, *Yo-Yo Ma played the cello*, are less likely to have Co-agents. Appropriate word sense disambiguation can improve the performance of semantic role systems by narrowing the sets of semantic labels that can apply.^[53]

A fundamental problem for WSD is choosing the set of senses to be distinguished. Generally referred to as a sense inventory, the set of senses used for a WSD system is expected to be a comprehensive and fixed list of senses for every word used in the domain of the application. This conception of word senses matches people's experience with dictionaries. Dictionaries encourage to consider words as having a discrete set of senses, yet any comparison between dictionaries quickly reveals how differently a word's meaning can be divided into separate senses. Rather than having a finite list of senses, many words seem to have senses that shade from one meaning into another; where to 'draw' the line between senses often seems to be an arbitrary decision. Moreover, the determination of how many lines to draw, that is, how narrow or how general to make the senses (this is an issue of **granularity**), can vary greatly, depending on who is creating the resource. WordNet is deemed to make **fine-grained** sense distinctions, whereas the sense distinctions made by PropBank and EngVallex are particularly **coarse-grained** and are only concerned with different senses that also have different subcategorization frames.

Different choices of sense distinctions can impact on semantic role frameworks. Fine-grained sense distinctions will encourage fine-grained semantic role labels in order to capture subtle differences in the types of argument that typically occur with these senses. Examples of shifting levels of granularity can be seen in the discussion of mappings between LIRICS and PropBank, and between VerbNet, EngVallex, and FrameNet.^{[2] [3]} With a few exceptions, VerbNet roles are at a level of granularity that is similar to that of LIRICS, while PropBank and EngVallex are typically more coarse-grained and FrameNet is typically more fine-grained.

6.1.4 Semantic classes

One view of lexical semantics postulates that word senses can be associated with concepts and, in particular, that different senses of a word are often characterized by different semantic classes associated with the syntactic object. How these associations are created and maintained is an open question. Katz and Fodor (1963)^[27] proposed one of the first linguistic theories of lexical semantics; it assumes that a linguistic grammar has a semantic component that assigns semantic representations to lexical items and, by means of recursive 'projection rules', to phrases and sentences. These representations are deemed to derive from a shared categorization of the world that includes common taxonomic elements such as PERSON, ANIMAL, PLANT, TREE, DOG, ARTEFACT, and BUILDING. This type of categorization relies heavily on subtype and supertype relations, such as *A DOG is a subtype of ANIMAL*, or *A TREE is a subtype of PLANT*.

Notwithstanding various attempts at building large-scale ontologies that capture these subtype relations in semantic class hierarchies, the field of natural language processing has never reached a consensus on a single, universal knowledge representation schema. However, the vast majority of systems rely on WordNet to fulfil this function, even though it was not developed with this in mind. The subtype/supertype relations mentioned in the previous paragraph parallel the lexical hyponym and hypernym relations that form the backbone of WordNet's associations. Other important lexical relations in WordNet are synonymy (similar, possibly substitutable, items), meronymy (part-whole), and antonymy (opposites). WordNet uses these relations extensively to create a large semantic network involving over 100 000 lexical units in the English language.^{[17] [38]} This network has proved to be an invaluable resource for natural language processing systems seeking semantic class inheritance relations, and these systems often use the hypernym relations in an ontological form to induce supertypes. This use of WordNet blurs the distinction between lexical items and purely conceptual categories. The approach has been ported to dozens of other languages.

Another example of a broad-scale natural language processing application of assigning semantic classes to data are the widely used Named Entity types, or nominal entity types, developed under the ACE (Automatic Content Extraction) program.^[51] This is basically a semantic classification task.

Nominal entity tagging is primarily focused on nouns and involves choosing a semantic category from a predefined category list (i.e. PERson, ORGAnization, GeoPoliticalEntity, LOCation, FACility, SUBstance, VEHicle, WEAPON) for each occurrence of the noun in context in a corpus. Several types of nouns, especially proper nouns such as ‘the White House’, can have multiple tags such as PER, GPE, and LOC. In these cases, determining which tag is appropriate, given a specific sentence as the context, amounts to the equivalent of a coarse-grained sense-tagging task (See 6.1.3). When annotators are engaged in nominal entity tagging, they typically work with a single set of entity classes that they choose from for each nominal they annotate. When they are sense-tagging polysemous items, each lexical unit has an individual set of senses and corresponding sense tags; it follows that when annotators move from one polysemous word to another in a sentence, they have a different set of senses to choose from for each one, and this makes sense tagging a more complex task.

Typical fillers for semantic roles are often referred to as being of a particular semantic class. Various applications have demonstrated that information about semantic class preferences can improve automatic semantic role labelling, especially when the roles being filled involve implicit arguments[23] [55], and it follows that this kind of information is helpful when associated with semantic roles.

6.1.5 Entailments

Different semantic roles are distinguished by the different inferences that they license. For example, in the sentence

(12) *Chris was drawing a picture on a sketchpad.*

the direct object, *a picture*, has, according to the LIRICS annotation scheme, the semantic role Result, which is defined as *participant in an event that comes into existence through the event. It indicates a terminal point for the event: when it is reached, then the event does not continue*. The sentence, therefore, licenses the following inferences about the picture that is referred to:

(13a) the picture did not exist prior to the drawing event;

(13b) the picture does exist after the drawing event.

Different senses of a verb are distinguished by differences in the semantic roles of their arguments; this is because they denote eventualities in which one or more of the participants are involved in a different way. They can also have different entailments; for example, in sentence (14),

(14) *Chris drew some water from a well.*

the verb *draw* is used in a different sense from that employed in sentence (12), and the direct object *some water* does not have the semantic role Result, but rather the role Theme; this is defined in the LIRICS annotation scheme as *“participant in a state or event that in the case of an event is essential to the event taking place but does not have control over the way the event occurs and is not structurally changed by the event.”* This sentence, therefore, licenses a number of inferences, including the following:

(15a) the water existed in the well prior to the drawing event;

(15b) there is also water outside the well after the drawing event.

Inferences such as those in examples (13) and (15) are called *entailments*. The concept of entailment is defined in mathematical logic as follows.

- *Entailment* is a relation that holds between a set of formulae: formula A and formula B. A entails B if each model (or ‘interpretation’ or ‘valuation’) of A is also a model of B (i.e. if B is true in every model of A). Expressed less formally, this means that formula B is true whenever the formulae in A are true.

The notion of entailment extends to natural language statements through the use of logic to represent meanings of natural language sentences. In recent years there has been much interest in assessing performance on textual entailment in the natural language processing community, as evidenced by a series of RTE (Recognizing Textual Entailment) tasks.[62]

6.2 Introduction to a metamodel for semantic role annotation

A predicate argument structure can be viewed semantically as formally representing an actual, generic, or hypothetical eventuality with its participants. Associated with the predicate (prototypically a verb) is an argument frame (typically called a 'subcategorization frame' for a verb) that describes the type of role performed by each participant expected to occur in that particular eventuality frame. Each slot in the eventuality frame can be given a semantic role label, which can then be associated with any participant that fills the slot. In the most fine-grained view, each individual lexical unit can be seen as defining a unique eventuality frame with a unique set of possible types of participant roles.

Different predicative expressions can share the same, or a very similar, set of possible participants. Obvious examples are nouns and adjectives that constitute derived forms of the same lexical item (*observe, observation, observer*). More complex examples involve distinct lexical items that can describe the same type of event from different perspectives, as in *buy* and *sell*, and *give* and *receive*. Depending on the desired level of generalization, the grouping of lexical items into shared eventuality frames may stop there (this is one view of the PropBank Frame Files), or it can continue to include a small set of items with very closely related semantics (the FrameNet view), or it may extend to include items that share specific patterns of argument types but can have a fairly tenuous semantic relation (the VerbNet view). These frameworks take the eventuality frame as a whole into consideration when determining the choice of individual semantic roles; this is motivated by examples such as *replace*, which can have one participant as the old item being replaced, and another participant as the new item replacing it, with an obvious dependency between these two roles. LIRICS does not presuppose the use of eventuality frames, but instead uses a set of features in the spirit of Dowty (1991)^[14] (e.g. like intentionality of the involvement of a participant) to distinguish between individual semantic roles and to decide which semantic roles to assign to the arguments of a predicate in a given utterance.

Note that, as discussed above, the same word can have multiple word senses, or lexical units, each of which may be associated with a distinct eventuality frame. In this case, the lemma can be represented by several eventuality frames, each one associated with a different frame or class. Therefore, for the approaches to semantic role labelling exemplified by FrameNet, PropBank, EngVallex and VerbNet, there are three core elements that shall be defined: the *word sense*, or *lexical unit*, under consideration; the *eventuality frame* associated with that word sense; and *specific semantic role labels* associated with each slot in that frame that will be assigned to the participants filling the slots. The more examples that can be provided to illustrate the degree of variation typical of the word in question, the better.

For FrameNet, the eventuality frames associated with a semantic role-labelling scheme are the FrameNet Frames which have multiple members; for PropBank, these are the PropBank rolesets or framesets associated with individual lemmas; for EngVallex, they are similar to PropBank Frame Files; and for VerbNet, they are associated with an entire class, so therefore with all the members of the class. The eventuality frames are typically consulted during the annotation process to guide the annotation and to ensure consistency. This makes the specification of the eventuality frames a critical step in the path towards the creation of an annotated corpus. For each predicate in a language, a meta-level description of the predicate and its arguments needs to be created, with examples; this description constitutes the definition of the eventuality frame. For a more detailed discussion of this process, see [C.2](#).

An additional consideration when defining any semantic role labelling scheme is the treatment of adjuncts, or eventuality modifiers. Schemes can differ in exactly which participants they would label as adjuncts and which as arguments, or they may finesse the question completely by giving unique labels to adjuncts for each eventuality frame.

Since annotations add linguistic information to spans of primary data, i.e. utterances, relevant spans in the data shall be identified. In stand-off format, this is realized either directly by means of pointers to the primary data (the original text or transcribed speech) or indirectly by referring to elements at another layer of annotation, such as a syntactic parse, of which the units identify regions of primary data. Following ISO practice, the term 'markable' is used to refer to the entities that 'anchor' an annotation in the primary data.

[Figure 1](#) visualizes the conceptual view that underlies semantic role annotation in this part of ISO 24617. Markables, anchoring semantic information to the primary data (a text segment) in a given source document, correspond to the linguistic realization of eventualities (actual or hypothetical states, events,

or facts) and their participants. Entities and eventualities can function as participants, and the mode of involvement of each participant in an eventuality is given a semantic role label with a specific semantic role type. Each semantic role relates one eventuality to one or more participants.

Note that virtually any kind of entity can be a participant in an eventuality: persons, physical objects, organizations, substances, locations, but also embedded eventualities, as in (16a); properties, as in (16b); amounts of something, as in (16c); numbers (16d); a method, as in (16e); a proposition, as in (16f); a set of entities, as in (16g), and so on.

(16a) *Eric and Nicole went to the concert.*

(16b) *He painted his house blue.*

(16c) *This weight of this suitcase exceeds 20 kgmes.*

(16d) *The birth of the twins increased the number of our children to four.*

(16e) *Let's decide by flipping a coin.*

(16f) *John believes that he will get the money together.*

(16g) *Marie, Jill and Chris got together.*

A predicative expression in natural language, in the sense in which it is understood in a given utterance, is viewed as denoting a certain type of eventuality, and its occurrence in the utterance as denoting an instance (or 'token') of that type of eventuality. Each eventuality frame specifies the semantic roles that the participants in such an eventuality may play. Semantic types of these semantic roles may also be specified. The same eventuality frame can be used to describe an action in progress or a completed action. This aspectual distinction is captured by the eventualityType, as illustrated by the examples in [Annex C](#). This is central to the instantiation of the Eventuality Frame.




Figure 1 — Metamodel for semantic role annotation

6.3 Criteria for distinguishing semantic roles

Each semantic role should, as far as possible,

- have a definition that characterizes the way in which a participant is involved in an eventuality, which could be operationalized by the specification of a set of licensed inferences (e.g. the participant acts intentionally or the participant exists independently of the eventuality),

- occur in corpus data in a substantial way,
- be consistently learnable by humans and machines,
- fit into an approach to semantic annotation with different levels of granularity,
- be generalizable across languages, genres, and domains,
- be relevant for obtaining good coverage of the ways in which participants can be involved in eventualities, and
- be plausibly mapped to one of the existing annotation schemes described in this part of ISO 24617.

6.4 Defining eventuality frames

Definition of the eventuality frames for a language is a lengthy, labour-intensive process that requires careful attention to detail and which benefits from a wide knowledge of lexical semantics, a close study of the relevant linguistics literature, and an examination of thousands of corpus instances. Many projects begin by selecting a large corpus that can provide examples for consideration when defining eventuality frames. This could be the corpus that is to be annotated, but if so, it is wise to supplement it with more general corpora. A balanced, broad-coverage corpus is more likely to facilitate a definition of eventuality frames that will generalize readily across domains, although it could be prohibitively labour-intensive to fully annotate such a corpus. The examples for each predicate are then examined to determine which predicate arguments typically occur with it, what semantic roles they play, and what the range of syntactic variation is. Distinctions between senses ('lexical units' in FrameNet parlance) of the same lemma are often based on differences in their semantic roles. Once an eventuality frame has been created and used for annotation, it should be evaluated against other pre-existing lexical resources and randomly selected additional examples, and revised as needed. The semantic roles used in this eventuality frame might have already been defined along with criteria for assignment; in this case, they and their criteria might now be subject to revision. As more and more frames are added, these role definitions are constantly re-examined, as are the criteria for distinguishing between arguments and adjuncts. As the definitions stabilize, they should be mapped to the ISO semantic roles provided in [Annex A](#), or to a resource that already has such a mapping, such as PropBank, VerbNet, or FrameNet.

Defining all possible eventuality frames for a single language is a daunting task, and there is no absolute measure of accuracy that can be applied. A certain element of subjectivity enters into any discussion of lexical semantics. Even so, the highest possible standards of consistency and explanatory power shall be maintained, and substantial efforts shall be made to achieve Quality Control. Adherence to comprehensive lexical schemas that have already been defined and have been widely used is strongly recommended. However, for any new language, there could be language-specific idiosyncrasies that require unique treatment. Where these diversions are necessary, they shall be clearly documented and, as far as possible, mapped to the already existing lexical schema that is the closest, with detailed definitions of discrepancies. A detailed mapping to ISO semantic roles, as exemplified by the VerbNet/LIRICS mapping, is one important element of a Quality Control process that can help ensure consistency. Another key element is a continuing cycle of **annotate/revise, annotate/revise, annotate/revise**.

During the annotation process, it is very helpful to constantly monitor inter-annotator agreement when highlighting both eventuality frame definitions that might be vague, incomplete, or inconsistent, and annotation guidelines that require clarification. Constant feedback from the annotators with respect to unexpected phenomena, difficulty in applying definitions, and new vocabulary is critical. In order to ensure consistency, it is advantageous to have a single annotator annotate all instances of a single sense of a predicate within a single task in a single sitting. Another way of improving annotator consistency, where double-blind annotation is being used, is to start with gold-standard or automatic parses (which predefine the constituent spans) or with markables that the labels are then associated with. Applying an already defined set of eventuality frames to a new corpus can be a useful 'sanity check' on their coverage and portability. Effective use of the annotated data for supervised machine-learning and later inclusion in a standard NLP application can also test the consistency and coherence of the annotation and the usefulness of the semantic role distinctions.

7 Interactions

7.1 Semantic roles and semantic types

Predicates differ enormously with regard to the range of semantic types and level of specificity/generality associated with the participants that can serve as their arguments. For example, *remember* accepts a wide range of nominal objects and sentential complements; conversely, and for all practical purposes, *devein*, *mint*, and *diagonalise* almost inevitably require *shrimp*, some sort of coin, and *matrices*, respectively, as their arguments. Labelling the semantic role Patient for each of these is not particularly informative and does not provide the prototypical semantic type. Historically, lexical resources are not expected to list either semantic roles or semantic types down to the level of *shrimp*, *coins*, and *matrices*, although that could benefit NLP systems. Recent efforts by Patrick Hanks to apply empirical techniques to determine fine-grained Lexical Sets that can be associated with individual verbs could provide useful information for NLP.^[26]

Linguists have traditionally referred to a rather limited set of semantic types such as inanimate/animate/sentient/human; these have clear grammatical relevance and have proved adequate when describing a wide variety of linguistic phenomena across many languages. Some lexical resources associate semantic types with certain semantic roles; for example, the Agents (animate, intentional, initiators) of verbs of communication are expected to be sentient (prototypically human). In this part of ISO 24617, it is not attempted to resolve the question of whether this kind of semantic type information is

- a) a constraint imposed on fillers of the Agent role by communication verbs,
- b) an entailment afforded by them, or
- c) simply a shorthand for describing the set of properties shared by most fillers of this role.

Failure to match expected constraints on the semantic type of a role is often used as a trigger for coercion.^[45]

7.2 Complexities

When dealing with more complex situations, the interplay between semantic roles becomes correspondingly more complex. Let us take commercial transactions as an example. These involve a minimum of four participants, which can be labelled Buyer, Seller, Goods, and Services, and two transfers of ownership, which we will call Goods_transfer and Money_transfer, and each transfer has an Agent/Source and a Recipient. When using the verb *buy*, the agency of the Buyer is profiled and the Seller and the money transfer are de-profiled, but when the verb *sell* is used, the converse is true. PropBank (EngVallex) labels the Agent Arg0 (ACT), the Theme (either Goods or Money), Arg1 (PAT), and the Recipient Arg2 (ADDR) for both transfers. FrameNet takes a different approach and labels the participants with the frame-specific semantic roles Buyer, Seller, Goods, and Money, regardless of whether they are Agents, Recipients, or Themes.

7.3 The role of context in semantic role assignment

Although it is often assumed that semantic roles are assigned solely by the lexical properties of the predicate, it appears that contextual factors also play a role in this process. A case in point involves verbs that allow optional agency in their subjects; this appears to be licensed by the inherent properties of the argument fillers (either the filler of the subject role or that of a different role) and/or other contextual cues such as adverbial modification, purposive constructions, and pragmatic inferences.

The following are examples from Van Valin and Wilkins (1996),^[69] and Wechsler (2005).^[70] The analysis of the examples varies according to the underlying assumptions about verb senses, the verb lexical representation, and the background knowledge at play.

(17a) *The thief broke the window.*

(17b) *The baby broke the window.*

(18a) *John hit the ball.*

(18b) *John hit a lamppost.*

(19a) *The sun disappeared behind the mountain.*

(19b) *John always disappears when work is mentioned.*

(20a) *John rolled down the hill.*

(20b) *John rolled down the hill as fast as he could in order to get to the road before the bikers got there.*

This calls for an approach to semantic role analysis that incorporates a rich characterization of contextual structure that is compatible with compositional processes.

7.4 Fuzzy boundaries between roles

It is sometimes quite difficult to decide on a semantic role label for a particular participant in an eventuality. The classic definition of a Patient is that it undergoes a fundamental change of state. Changes of location are not usually considered to be fundamental enough to warrant a Patient label, so items in motion are typically labelled Themes instead of Patients. However, if someone is *tarred and feathered and ridden out of town on a rail*, the question arises as to whether or not tarring and feathering would so affect the victim that they would be labelled a Patient rather than a Theme. Annotators have to make many subtle judgements such as this when faced with individual sentences, and the distinction between Patient and Theme can vary depending on the annotator's subjective interpretation of how affected a participant is by the eventuality. Distinctions between Agent and Cause or Agent and Stimulus might depend on an annotator's subjective interpretation of how intentional a participant's actions are; for example, annotators disagree as to whether entities that are not affected in their physical integrity but undergo termination of their functionality as a result of the event should be labelled Patient.

7.5 Multiple classification

Annotators also have to make subtle judgements when they annotate contexts where one participant appears to share entailments belonging to different roles. For example, in (21), the subject can be said to share entailments belonging to both the Agent role and the Experiencer role. In such cases, LIRICS allows the assignment of multiple semantic roles to a participant.

(21) *The tourists admired the painting.*

7.6 Inheritance relations between semantic roles

Given that different semantic role resources might choose to define semantic roles at different levels of granularity, inheritance relations naturally arise as a vehicle for mapping between the different resources. A lot of traditional research on semantic roles assumes a hierarchy of semantic roles with strong entailment relations, but such an approach might not be appropriate in all semantic domains in all languages. For illustrative purposes, this subclause describes a thematic role hierarchy that has been developed for VerbNet and was inspired by the mapping between VerbNet and LIRICS (not all systems use an explicit hierarchy, e.g. EngVallex uses only a flat list of semantic roles). It offers fine-grained roles that can highlight subtle differences in the behaviour of individual verb classes, coarse-grained thematic roles that can be used to describe the most general aspects of verbal behaviour, and an intermediate level that is maximally descriptive of verbal behaviour while also generalizable beyond specific verb classes. With this hierarchy, users can select the level of granularity that suits their task best. A coarse-grained superordinate role can be used in place of a finer-grained role for tasks that require a roleset, which in turn has the broadest coverage across all verbs; conversely, fine-grained and class-specific roles can be used for tasks that benefit from information that helps to differentiate between classes of verbs. The VerbNet thematic role hierarchy is given in [Figure 2](#). Further description and examples illustrating several hierarchy components follow.




Figure 2 — The VerbNet thematic role hierarchy

The hierarchy begins with a root node of Participant^[63], which dominates the proto-roles Actor and Undergoer^[64]. Other meta-roles immediately below Participant include Place and Time. These roles are all defined very loosely, and Place roles are therefore not strictly spatial locations, and can include abstract directional roles as well as Paths. Undergoer could be extended to include Manner and Attribute. Roles that are not in a parent-child relationship can co-occur; for instance, Agent and Patient co-occur frequently, but Patient and Experiencer (a child of Patient) never co-occur. Parent-child relationships are governed by additional restrictions placed on the lower node, and in principle, nodes that include restrictions can be thought of either as named semantic roles, such as Recipient, or as higher-level parent nodes plus feature sets (for Recipient, Goal [+concrete][+animate]).

The highest levels of this hierarchy, or the superordinate roles, are applicable to all verbs. As has been previously made clear, the addition of restrictions to these roles forms the basis of subordinate roles. For example, in the hierarchy, Experiencer is subordinate to Patient, and therefore inherits all features of Patient but is characterized by the additional restriction *+awareness*; accordingly, an Experiencer is a Patient who is aware of the process denoted by the verb that the participant is undergoing and is affected by it. The role Patient is underspecified for awareness. In some cases, the subordinate roles will be specific to certain event classes; for example, Topic is a Theme with the additional restriction *+information_content*, and a Topic is therefore a Theme in events of information transfer or communication verbs. Examining the PropBank annotations reveals that Topics are generally realized syntactically as “that” clauses or eventualities, e.g. ‘She said that she would go’, while Themes tend to be entities realized as noun phrases. Keeping the roles distinct can facilitate automatic classification based on syntactic constituent types.

The hierarchy provides the user with flexibility in using the resource, and this flexibility also facilitates mapping between other semantic resources, as described in [Annex B](#).

8 Guidelines for developing new semantic role frameworks for languages and/or domains

8.1 General

A recommended and speedy approach to creating an interoperable semantic role labelling scheme is to start with a framework that has already been defined, and to show how it can be applied to a new

language or domain. It is likely that certain modifications will be needed to fit this scheme appropriately to the new application area, and documentation of these modifications and of any other variations is crucial for ensuring interoperability. The next subclause, 8.2, describes the results of mapping between LIRICS and VerbNet semantic role definitions, and is an example of how this could be presented.

8.2 Mapping from VerbNet to LIRICS

A systematic comparison of VerbNet (VN) with another semantic resource, LIRICS (Linguistic Infrastructure for Interoperable Resources and Systems, described in detail in [Annex A](#)), reveals the differences in granularity between the semantic roles of the two resources, and demonstrates the utility of a pivotal set of ISO semantic roles. The LIRICS semantic roleset was created with the aim of potentially serving as an ISO standard set of semantic roles, and was therefore intended to be general enough to be practical for a wide variety of annotation and event types. VerbNet was created as a resource for a variety of NLP tasks and is organized into verb classes; the semantic roles in VerbNet differ in several key aspects from those of LIRICS. The comparison between LIRICS and VerbNet roles is provided in [Table 1](#). Notice that VerbNet has roles that do not have a mapping to a LIRICS role, as well as several roles that map to a single LIRICS role. This sometimes reflects the greater specificity of VerbNet roles: these can be unique to certain event types, and they are therefore not included in the LIRICS set of roles.^[2]

Table 1 — Original mapping from VerbNet semantic roles to LIRICS

VNRole	LIRICS Role	VNRoles ctd.	LIRICS Roles ctd.
Actor	Agent	Patient	Patient
Actor1	Agent	Patient1	Pivot
Actor2	Partner	Patient2	Patient
Agent	Agent	Predicate	—
Asset	Amount	Product	Result
Attribute	Attribute	Proposition	—
Beneficiary	Beneficiary	Recipient	Goal
Cause	Cause, Reason	Source	Source, Initial_Location
Destination	Final_Location	Stimulus	Theme
Experiencer	Pivot, Patient	Theme	Theme
Extent	Amount, Distance	Theme1	Theme
Instrument	Instrument	Theme2	Theme

The ideal set of semantic roles should be able to label the arguments of any relation concisely; however, what this set of roles should be has long been a subject of dispute in the linguistic community. Some of the role choices in VerbNet reflect differences in the syntactic patterns of coarse-grained and fine-grained roles; for example, the fine-grained role Topic, which is specific to verbs of communication, is more likely to be realized in the form of a complement clause, unlike the more coarse-grained role Theme, which is more likely to be realized in the form of a noun phrase. LIRICS does not draw any distinctions related to syntactic structure. Defining the mapping to the LIRICS roles was beneficial to VerbNet in many ways, since it encouraged discussion on, and clarification of, the semantic role definitions, and ultimately led to a simplification of the framework. Actor, Actor1, Actor2 were dropped completely, and all of the Theme1, Theme2, Patient1 and Patient2 roles were also dropped in favour of the Co-Theme and Co-Patient for the second mention of the related role. The LIRICS Partner role was seriously considered, however, given the possibility of plural Agents and plural Patients occurring in the same sentence, there was a concern Partner could be ambiguous. The LIRICS Pivot role was adopted for verbs of possession. It was eventually decided to keep Experiencer/Stimulus, but they are now restricted to a smaller set of verbs where more narrowly defined definitions of the roles clearly apply. The new mapping that results from these changes is set out below in [Table 2](#). A few minor differences remain: they include different uses of Extent and Amount, and differences in granularity represented by different names for Destination and Final_Location, Product and Result, and a few others.^[3]

Table 2 — Revised mapping between VerbNet and LIRICS

VerbNetRole	LIRICS Role	VerbNetRoles ctd.	LIRICS Roles ctd.
Agent	Agent	Theme	Theme, Pivot
Co-Agent	Partner	Co-Theme	Partner
Asset	Amount	Topic	Theme
Attribute	Attribute	Product	Result
Beneficiary	Beneficiary	Predicate	—
Cause	Cause, Reason	Recipient	Goal
Destination	Final_Location	Source	Source, Initial_Location
Experiencer	Pivot, Patient	Stimulus	Theme
Extent	Amount, Distance	Patient	Patient
Instrument	Instrument	Co-Patient	Partner

Key to integrating an ISO standard set of semantic roles within the organization of an application-oriented semantic role framework like VerbNet is a mapping to the standard that will make it possible to take advantage of the strengths of a set of roles applicable to all verbs, while maintaining the practicalities associated with roles that help to distinguish certain classes of verb. A semantic role hierarchy, like that of VerbNet (described in B.3.2), can help to facilitate mapping across resources of differing granularity, because more fine-grained roles can be used when mapping to resources of greater semantic specificity, such as FrameNet, while more coarse-grained roles can be used when mapping to resources that are more semantically general. The examples set out below show how the current VerbNet thematic role hierarchy allows different levels to map more appropriately to different resources. Notably, the fine-grained levels of the VerbNet thematic roleset largely overlap with FrameNet, while the more coarse-grained levels of the hierarchy largely overlap with LIRICS.

- a) He talked about politics.
- FrameNet (Statement frame): He_{SPEAKER} talked_{RELATION} about_politics_{TOPIC}
 - VerbNet Fine-Grained (Talk-37.5): He_{AGENT} talked_{RELATION} about_politics_{TOPIC}
 - VerbNet Coarse-Grained (Talk-37.5): He_{AGENT} talked_{RELATION} about_politics_{THEME}
 - LIRICS: He_{AGENT} talked_{RELATION} about_politics_{THEME}
- b) He sent the letter to Mary.
- FrameNet (Sending frame): He_{SENDER} sent_{RELATION} the_letter_{THEME} to_Mary_{RECIPIENT}
 - VerbNet Fine-Grained (Send-11.1): He_{AGENT} sent_{RELATION} the_letter_{THEME} to_Mary_{RECIPIENT}
 - VerbNet Coarse-Grained (Send-11.1): He_{AGENT} sent_{RELATION} the_letter_{THEME} to_Mary_{GOAL}
 - LIRICS: He_{AGENT} sent_{RELATION} the_letter_{THEME} to_Mary_{GOAL}
- c) The contractor builds houses.
- FrameNet (Building): The_contractor_{AGENT} builds_{RELATION} houses_{CREATED_ENTITY}
 - VerbNet Fine-Grained (Build-26.1): The_contractor_{AGENT} builds_{RELATION} houses_{PRODUCT}
 - VerbNet Coarse-Grained (Build-26.1): The_contractor_{AGENT} builds_{RELATION} houses_{RESULT}
 - LIRICS: The_contractor_{AGENT} builds_{RELATION} houses_{RESULT}

Annex A (informative)

Specification of ISO semantic roles

A.1 LIRICS

The LIRICS description model and semantic roleset incorporate important findings from other projects, including FrameNet, PropBank, and VerbNet, in the same area, and take a step forward by providing a complete set of semantic roles without redundancies, defined as purely semantic concepts by virtue of a set of distinctive semantic properties. The LIRICS model encompasses different levels of granularity that enable hierarchical structures of semantic roles and make this model extendable and attractive for many applications. It has been established that the LIRICS semantic roleset can be used reliably for annotation purposes, and annotators demonstrate substantial agreement using the LIRICS data categories.

The aim of the EU-funded LIRICS project, which was set up as a spin-off of ISO TC 37/SC 4, is to explore the possibility of establishing sets of annotation concepts, defined in accordance with ISO 12620 as so-called data categories, for syntactic, morpho-syntactic, and semantic annotation, and lexical markup. In the section of the project that addresses semantic annotation, several approaches and existing annotation schemes for semantic role labelling are analysed and compared as to

- description model,
- granularity,
- definition of semantic roles, and
- consistency and reliability of annotation.

On the basis of 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,
- not as primitives but as relational concepts that link participants to events, and
- as generally applicable rather than restricted to only a few specific verb, noun, or adjective classes.

LIRICS specifically defines semantic roles as relational notions that link a participant to some real or imagined event, state or fact ('event'), and describes the way the participant is involved, rather than a participant's internal properties (such as sentience). Starting with the most frequently used roles (e.g. *Agent*, *Theme*, and *Patient*), a list of entailments was established for each role; they apply to a participant in that role, regardless of the type of event. The boundaries between roles were examined with a view to designing a maximally complete and minimally redundant roleset, and these entailments were converted into a set of properties, for example, [$\pm$ intentionality], [$\pm$ independent existence].

Table A.1 — Semantic properties for *Theme* and *Result* roles

Theme	Result
- intentionality	- intentionality
- affectedness	- affectedness
+ independent existence	- independent existence

Table A.1 illustrates the differences between *Theme* and *Result* in terms of these properties, and this is how a set of 29 ‘high-level’ roles was constructed.[34] [35]

NOTE LIRICS defines 11 roles which are central to any event (e.g. Agent, Theme, and Patient), 10 adjunct roles (e.g. Time, Location, and Manner) and eight sub-roles for Time and Location (e.g. Duration, Frequency, and Path). For definitions and illustrative examples of each individual semantic role, see References [34] and [35].

Table A.2 lists the definitions of the LIRICS semantic roles in the form of ISO data categories.

Table A.2 — LIRICS semantic roles, defined in the format of ISO data categories

agent	
Definition	Participant in an event who intentionally or consciously initiates an event, and who exists independently of the event.
– Source	Adapted from Dowty [1989],[71] EAGLES, SIL, Sowa [2000],[72] and UNL
Explanation	An agent may be animate, or only seemingly so, or perceived as animate; this is to ensure that cases of nonhuman agency such as a robot or an institution will not be excluded from being able to initiate an event (e.g. ‘GM offers rebates on its new models’).
Example	‘John [agent e1] built e1 the house’
beneficiary	
Definition	Participant in an eventuality that is advantaged or disadvantaged by the eventuality, and that exists independently of the event.
– Source	Adapted from EAGLES, Sowa [2000],[72] and UNL
Explanation	The roles of beneficiary and recipient differ in that the recipient is the ultimate target of an action, whereas a beneficiary is not. They also differ in that the role of recipient is the ultimate goal of an action, whereas that of beneficiary is not. Furthermore, a recipient cannot take part in a state.
Example	‘John sold e1 the car (for a friend [beneficiary e1])’ ‘He gave e1 his life for his country [beneficiary e1]’
cause	
Definition	Participant in an event that initiates the event, but does not act with any intentionality or consciousness; the participant exists independently of the event.
– Source	Adapted from SIL (‘Causer’) and Sowa [2000][72] (‘Effector’)
Explanation	Except for the lack of intentionality of the participant, this semantic role is very similar to that of the agent, and in fact shares all its other properties. The role of cause can often be identified with verbs of initiation or causation such as ‘to cause’, ‘to produce’, ‘to start’, ‘to originate’, ‘to occasion’, and ‘to generate’.
Example	‘The wind [cause e1] broke e1 the window’ ‘His talk [cause e1] produced e1 a violent reaction e2 from the crowd.’
goal	
Definition	Participant in an event that is the (non-locative, non-temporal) end point of an action; the participant exists independently of the event.
– Source	Adapted from Sowa [2000][72]
Explanation	Goal differs from beneficiary in that there need not be a clear benefit or disadvantage for the participant; moreover, only <i>events</i> can have a participant towards whom the action is directed.
Example	‘(The executive [agent e1]) recalls e1 ((Mr Corry [agent e2] whispering e2 (to him and others [goal e2]) [theme e1]).’
instrument	
Definition	Participant in an event that is manipulated by an agent, and with which an intentional act is performed; it exists independently of the event.

Table A.2 (continued)

– Source	Adapted from EAGLES ('Implement'), SIL, Sowa [2000],[72] and UNL
Explanation	Some [Loos et al. 2004][73] would define instrument more simply as an inanimate object used to implement an event. However, it is clear that the semantic role of instrument can also be animate (e.g. 'The woman was dragged by her horse [instrument] for several meters', or even 'John [instrument] threw himself at the door in a rage'), and can also take part in a state as well as an event (e.g. 'The tarpaulin is tied down with rope [instrument]').
Example	'He opened e1 the door with the key [instrument e1] '(The brick [instrument e1] [cause e2]) hit e1 the window and shattered e2 it'
partner	
Definition	Participant in an event that is intentionally or consciously involved in carrying out the event. Participant is not the principal agent of the event and exists independently of the event.
– Source	Adapted from FrameNet, PropBank, and UNL
Explanation	The semantic role of partner mainly differs from that of agent in that the participant in question is performing the action in accompaniment to the agent, and is not the primary focus. Sometimes partner is also known as 'accompaniment' (SIL and Sowa, 2000).[72] UNL distinguishes the roles co-agent and partner, but this distinction is felt to be unclear.
Example	(Libyan leader Muammar Gaddafi [agent e1]) publicly [manner e1] announced e1 (his commitment e2) (to dismantle e3) (WMD programs [patient e3]) (in his country) [location e3] [theme e1] [purpose e2]) following (negotiations e4 (with US and UK authorities [partner e4]) [reason e1])
patient	
Definition	Participant in an event that undergoes a change of state, location, or condition, is causally involved or directly affected by other participants, and exists independently of the event.
– Source	Adapted from EAGLES, Sowa [2000],[72] and UNL
Explanation	Patient is distinguished from the semantic role of theme by the fact that it is structurally changed or affected by the event.
Example	(White women [agent e1] [agent e2]) serve [e1] (tea and coffee [theme e1]), and then wash [e2] (the cups and saucers [patient e2]) (afterwards [time e2])
pivot	
Definition	Participant in a state that is characterized as being in a certain position or condition throughout that state, and has a major or central role or effect in that state. A pivot is more central to the state than a participant in a theme role and exists independently of the state.
– Source	Webster New Collegiate Dictionary [1976]
Example	(Vicar Marshall [agent e1, pivot e2]) admits [e1] (to mixed feelings [e2] (about this issue [theme e2]) [theme e1])
purpose	
Definition	Set of facts or circumstances that an agent wishes or intends to accomplish by performing some intentional action.
– Source	Adapted from EAGLES and UNL
Explanation	The notion of purpose implies intentional action. This property of purpose makes it impossible to have this semantic role in a state. The role of purpose differs from that of reason in that purpose describes the aims of an agent, whereas reason indicates why the event is carried out or the state is true. Contrast the following examples: (a) 'The authorities extradited him to the US for drug trafficking [reason]' and (b) 'The authorities extradited him for trial [purpose] in the US'.
Example	'(The robber [agent e1]) (tied e1) Harry [theme e1] (to the chair [final-location e1]) ((to stop e2) him [theme e2] (from getting away [purpose e2]) [purpose e1])'
reason	
Definition	Set of facts or circumstances explaining why a state exists or an event occurs.

Table A.2 (continued)

– Source	Adapted from Sowa [2000][72] ('Matter') and UNL
Explanation	The role of reason can be distinguished from that of purpose by the fact that the latter indicates the objective or goal of an agent that acts intentionally. However, the role of agent is not applicable to states. Reason is also different from manner and method in that reason describes why the event is being carried out, while manner and method describe how it is being carried out.
Example	'People love giant pandas because they look like teddy bears [reason]' 'Due to the spray from the hose [reason] the ink ran down the page before she could read it.'
result	
Definition	Participant in an event that comes into existence through the event; it indicates a terminal point for the event: when that is reached, the event does not continue.
– Source	Adapted from Sowa [2000][72]
Explanation	Result is the completed point of a process and, unlike goal, is dependent on the event for its existence.
Example	'(Within the past two months [duration e1]) (a bomb [cause e1]) exploded e1 (in the offices of El Espectador in Bogotá [location e1]), (destroying e2 (a major part of its installations and equipment [patient e2]) [result e1])'
setting	
Definition	Set of (non-locative and non-temporal) facts or circumstances of the occurrence of an event or a state.
– Source	Adapted from EAGLES ('Scene'), UNL ('Condition'), and FrameNet ('Circumstances')
Example	'(A number of medical and agricultural research centers [pivot s1][instrument e1]) had s1 (the potential [attribute s1]) to be used e1 (in BW research [setting e1]).'
source	
Definition	Non-locative, non-temporal starting point of an event. The source exists independently of the event.
– Source	Adapted from Sowa [2000]
Example	'(Eaton [beneficiary e1]) earned e1 (from continuing operations [source e1])'
theme	
Definition	Participant in a state or an event that (a) in the case of an event, is essential to the event taking place, but does not have control over the way the event occurs and is not structurally changed by the event, and (b) in the case of a state, is characterized as being in a certain position or condition throughout the state, and is essential to the state being in effect but not as central to the state as a participant in a pivot role. The theme of a state or event exists independently of the state or event.
– Source	Adapted from EAGLES, Sowa [2000], and UNL ('Object')
Explanation	The Theme role is distinguished from the Patient role by the fact that it is not structurally changed by the event or in the state in which it occurs.
Example	'(One man [agent e1]) wrapped e1 (several diamonds [theme e1]) (in the knot of his tie [final-location e1])'
time	
Definition	Participant that indicates an instant or a time interval during which a state exists or an event takes place.
– Source	Adapted from EAGLES, SIL, Sowa [2000], and UNL
Explanation	As with location, time is divided into three subroles: beginning, end, and duration.
Example	'(Right now [time e1]) ((about a dozen [amount e1]) standardiseatories [agent e1][agent e2], (in the US. Canada and Britain [location e1]), (are racing [e1]) (to unmask e2 (other suspected tumour-suppressing genes [theme e2]) [purpose e1])'

Table A.2 (continued)

manner	
Definition	The way or style of performing an action or the degree/strength of a cognitive or emotional state.
– Source	Adapted from EAGLES and UNL
Explanation	The role of manner differs from that of instrument in that the former describes an event as a whole, whereas the latter characterizes one of the components of the event or state. Manner is abstract, whereas instrument is concrete. The latter also differentiates instrument from means and method. The role of manner includes secondary effects (quietly, loudly), and general descriptions comparing events or states (in the same way). It may also indicate salient characteristics of theme, experiencer, agent etc. (e.g. coldly, deliberately, eagerly, and carefully).
Example	‘The tiny stick (was fastened s1) tightly [manner s1] to his wrist.’
medium	
Definition	The physical setting, device, or channel that allows an event to take place.
– Source	Adapted from Sowa [2000]
Explanation	The semantic role of medium and that of instrument are similar. The main difference is that an instrument is intentionally used by an agent, whereas a medium can occur in a non-intentional event, as the example given below illustrates.
Example	‘The students accidentally heard e1 the news on (the radio [medium e1])’
means	
Definition	Procedure for performing an action in terms of component steps, or a methodology by which an intentional act is performed by an agent. A means does not necessarily exist independently of the event.
– Source	Adapted from UNL
Explanation	Means differs from instrument in that the former describes possible activities and methods for doing something that may have no existence independent of the event, while the latter describes previously existing tools. This distinction is exemplified by (a) ‘I sliced the cucumber in 1/8th inch slices with a knife [instrument]’ and (b) ‘I sliced the cucumber in 1/8th inch slices by marking the intervals with a ruler [means]’.
Example	‘The mayor delayed e1 the ribbon ceremony (by pretending to be ill [means e1])’ ‘He had to button e1 his sleeve (by holding e2 the cuff in his mouth [means e1])’
location	
Definition	Place where an event occurs, or a state is true, or a thing exists.
– Source	Adapted from EAGLES (‘Place’), SIL (‘Locative’), and Sowa [2000]
Explanation	
Example	‘Here [location e1] is s1 (an example [theme s1])’ ‘She was cooking e1 (in the kitchen [location e1])’
initial-location	
Definition	Participant in an event that indicates the location where an event begins or a state becomes true; initial-location exists independently of the event.
– Source	Adapted from EAGLES, Sowa [2000] (‘origin’), and UNL (‘Initial Place’)
Example	‘(Half way out of the harbour [initial-location s1]) the sea becomes s1 really deep’
final-location	
Definition	Location where an event ends or a state becomes false; final-location exists independently of the event.
– Source	Adapted from EAGLES, Sowa [2000] (‘Destination’), and UNL (‘Final Place’)

Table A.2 (continued)

Example	'(One man [agent e1]) wrapped e1 (several diamonds [theme e1]) (in the knot of his tie [final-location e1])'
distance	
Definition	Length or extent of space that plays a role in an eventuality.
– Source	Adapted from WordNet
Example	'missiles [pivot s1] (capable s1) of (travelling e1) (more than 300 km [distance e1])'
duration	
Definition	Length or extent of time during which an event occurs or a state is true.
– Source	Adapted from WordNet
Example	'Terry [agent e1] jogged e1 (for two hours [duration e1])' 'I flew e1 to Hong Kong (in less than 10 h [duration e1])'
initial-time	
Definition	Indication of the point in time when an event begins or a state becomes true.
– Source	Adapted from EAGLES, Sowa [2000], and UNL
Example	'(Harry [agent e1]) (teaches e1) (on Friday (from 14:45 [initial-time e1]) (to 12:30 [final-time e1]) [time e1])'
final-time	
Definition	Indication of a point in time when an event ends or a state ceases to be true.
– Source	Adapted from EAGLES, Sowa [2000], and UNL
Example	'(Harry [agent e1]) (teaches e1) (on Friday (from 14:45 [initial-time e1]) (to 12:30 [final-time e1]) [time e1])'
path	
Definition	Intermediate location or trajectory between two locations, or in a designated space, where an event occurs.
– Source	Adapted from Sowa [2000]
Example	'(The baby [agent e1]) crawled e1 (across the floor [path e1])'
amount	
Definition	Quantity of something other than time or space, or number of objects of a certain kind, which plays a role in an event or a state.
– Source	Adapted from PropBank ('Extent'), FrameNet ('Amount'), UNL ('Quantity'), Sowa [2000] ('Amount', 'Measure') and EAGLES ('Quantity')
Example	'(The euro [theme s1]) (is worth s1) (nearly one and a half dollars [amount s1])' '(A flight to Hong Kong [theme e1]) takes e1 (about 12 h [amount e1])' 'When the twins were born, (the number of kids [theme e1]) increased e1 (by two [amount e1])'
attribute	
Definition	Property that an event or state associates with one of the other participants.
– Source	Adapted from FrameNet ('Parameter'), UNL, Sowa [2000], and EAGLES
Example	'(We [agent e1]) will (paint e1) (the front door [theme e1]) (dark green [attribute e1])' Dutch: '(Zij [agent e1]) (stampte e1) ([de korrels] patient e1) (fijn [attribute e1])'

A.2 LIRICS related to VerbNet, PropBank, FrameNet, and EngVallex

Table A.3 relates the LIRICS semantic roles to roles defined in VerbNet, PropBank, FrameNet (comparison between rolesets made possible due to the analyses made in the SemLink project.^[36]), and EngVallex

(=FGD/PDT). The table shows that the role of the LIRICS roleset is clearly comparable to that of VerbNet. However, upon examining the definitions, it is seen that VerbNet's roles are not purely semantic: they are partly defined as syntactic or lexical structures, and the semantic differences between roles are not captured. For example, VerbNet defines *Agent* as: '*generally a human or an animate subject, used mostly as a volitional agent, but also used in VerbNet for internally controlled subject such as forces and machines*'. Using 'subject' here, in the sense of grammatical subject, means that the definition does not apply to passive constructions as in *The horse was hit by the tractor*, where *the tractor* should be assigned an *Agent* role.

Table A.3 — LIRICS semantics roles related to roles defined in VerbNet, PropBank, FrameNet, and EngVallex (=FGD/PDT)

VerbNet	PropBank	FrameNet	LIRICS	FGD/PDT
Agent	Arg0, Arg1	Agent, Speaker, Cognizer, Communicator, Ingestor, Deformer etc.	Agent	ACTor
Actor	Arg0	Avenger, Communicator, Item, Participants, Partners, Wrongdoer	Agent	ACTor ADDRessee PATient
Actor1	Arg0	Arguer1, Avenger, Communicator, Interlocutor1, Participant 1 etc.	Agent	ACTor
Actor2	Arg1, Arg2	Addressee, Arguer2, Injured Party, Participant2, Partner2	Partner	ADDRessee PATient
Attribute	Arg1, Arg2	Attribute, Dimension, Extent, Feature etc.	Attribute	ACMP EXTent
Beneficiary	Arg1, Arg2, Arg3, Arg4	Audience, Beneficiary, Benefitted party, Goal, Purpose, Reason, Studio	Beneficiary	ADDRessee BENefactive EFFect
Cause	Arg0, Arg1, Arg2, Arg3	Addressee, Agent, Cause, Communicator etc.	Cause, Reason	ORIGin CAUS
Destination	Arg1, Arg2, Arg5	Addressee, Body part, Context, Goal etc.	Final Location	ACTor ADDRessee PATient AIM
Experiencer	Arg0, Arg1	Cognizer, Experiencer, Perceiver etc.	Pivot, Patient	ACTor
Extent	Arg2	Difference, Size change	Amount, Distance	EXTent DIFFerence
Instrument	Arg2	Agent, Fastener, Heating instrument, Hot Cold source etc.	Instrument	MEANS
Location	Arg1, Arg2, Arg3, Arg4, Arg5	Action, Area, Fixed location etc.	Location	LOC DIR1 DIR2 DIR3
Material	Arg1, Arg2, Arg3	Components, Ingredients, Initial entity, Original, Resource, Undergoer	Source	ORIGin
Patient	Arg0, Arg1, Arg2	Addressee, Affliction, Dryee, Employee, Entity, Executed etc.	Patient	ADDRessee PATient EFFect
Patient1	Arg0, Arg1	Concept 1, Connector, Fastener, Item, Item 1, Part 1, Whole patient	Pivot	PATient

Table A.3 (continued)

VerbNet	PropBank	FrameNet	LIRICS	FGD/PDT
Patient2	Arg2, Arg3	Concept 2, Containing object, Item 2, Part 2	Patient	PATient
Predicate	Arg1, Arg2	Action, Category, Containing event etc.	—	PREDicate
Product	Arg1, Arg2, Arg4	Category, Copy, Created entity etc.	Result	EFFect
Proposition	Arg1, Arg2	Act, Action, Assailant, Attribute etc.	—	ACTor PATient
Recipient	Arg1, Arg2, Arg3	Addressee, Audience, Authorities, Recipient	Goal	ADDRessee BENefactive
Stimulus	Arg1	Emotion, Emotional state, Phenomenon, Text	Theme	ACTor
Theme	Arg0, Arg1, Arg2	Accused, Action, Co-participant, Co-resident, Content, Cotheme etc.	Theme	ACTor ADDRessee PATient
Theme1	Arg0, Arg1	Cause, Container, Phenomenon 1, Profiled item, Theme	Pivot	PATient ORIGin EFFect
Theme2	Arg1, Arg2, Arg3	Containing object, Contents, Cotheme etc.	Theme	PATient ORIGin EFFect
Time	ArgM TMP	Time	Time	TWHEN TFHL THL THO TPAR TSIN TTILL
Topic	Arg1, Arg2	Act, Behaviour, Communication, Content etc.	Theme	PATient
Asset	Arg1, Arg3	Asset, Category, Measurement, Result, Value	Amount	PATient EFFect
Value	Arg1	Measurement, Result, Value, Asset, Category	Amount	PATient EFFect
Source	Arg2, Arg3	Role, Victim, Patient, Source, Path start etc.	Initial location	PATient ORIGin DIR1
—	—	Setting, ContainingEvent	Setting	
—	—	Means	Means	MEANS
—	ArgM Manner	Manner	Manner	MANNER
—	ArgM Purpose	Purpose	Purpose	Aim

The LIRICS set of semantic roles was evaluated with respect to redundancy, completeness, and reliability. Redundancy was tested by

- a) inspecting annotated data for the boundaries between semantic roles, aiming to avoid overlaps, and

- b) analysing the roleset using the defined set of properties and eliminating roles that are not clearly distinct in these terms. This led to the removal of a few roles such as *Stimulus* and *Experiencer*. *Stimulus* overlaps with *Theme*, and *Experiencer* either with *Patient* in an event or with *Pivot* in a state. The latter roles are broader concepts and, unlike *Stimulus* and *Experiencer*, are not restricted to mental, psychological, or perception events/states.

The completeness of the defined set of roles was measured both theoretically by comparing observations with the semantic rolesets defined in various other projects,^[42] and empirically. For the empirical evaluation of completeness and reliability, a multilingual test suite was constructed including English, Dutch, Italian, and Spanish. FrameNet and PropBank data were used for English. Three texts (120 sentences) and 83 isolated sentences were selected from the FrameNet corpus, and 355 sentences from the PropBank data. For Dutch, 15 texts were selected from news articles, with a total of 260 sentences. News articles were also used to construct the Italian part of the test suite (101 sentences), and all were taken from the Italian Treebank corpus. For Spanish, the test suite contained 189 sentences taken from the Spanish FrameNet corpus.

It was found that, in the material taken from FrameNet or PropBank data, the semantic roles as marked up with FrameNet or PropBank tags could be reliably re-annotated using the LIRICS roleset. The material that had not previously been annotated could also be completely tagged with the LIRICS roles, and the LIRICS roleset can therefore be considered to be (relatively) complete. (For completeness estimations comparing other projects, Petukhova et al. (2007)^[42] is referred here.)

The usability and reliability of the defined tagset were investigated by measuring inter-annotator agreement using the standard Kappa statistic. It was found that annotators reached *substantial agreement* (scores between 0,61 and 0,8) to *perfect agreement* (between 0,81 and 1,00) annotating semantic roles; the exceptions were *Instrument*, which was often confused with *Means* and with *Medium*, and *Source*, which was sometimes confused with *Reason*. (The definitions of these roles have since been formulated more sharply, and the annotation guidelines improved to avoid these confusions.) It was also found that spatial and temporal roles (*Location* and *Time*, and their sub-roles) are easier to identify than others.

Some situations, such as *Reason* vs. *Purpose*, are ambiguous:

(A.1) *Laws exist to prevent crimes.*

In this particular case, it is not entirely clear without context whether ‘preventing crimes’ is a *Reason* for ‘the existence of laws’ or a *Purpose*.

Since LIRICS defines semantic roles as a way in which a participant takes part in an event, and a participant’s involvement is potentially manifold, a participant can have more than one semantic role associated with an event. For example, for verbs like ‘pay’, ‘supply’, and ‘provide’, a participant who receives something can have two roles: not only *Beneficiary*, but also *Goal*:

(A.2) *Germany and China allegedly provided technical and material assistance to the Al-Fatah program.*

The participant *the Al-Fatah program* is clearly advantaged by the event (*Beneficiary*); it also forms a terminal point for the event (*Goal*).

Annex B (informative)

Review of existing frameworks

B.1 Overview

There are currently five English lexical resources that provide explicit semantic role labels for use in data annotation: FrameNet, VerbNet, LIRICS, EngVallex, and PropBank. (Note that NomBank, as a companion to PropBank, provides corresponding semantic role labels for noun predicates.^[37]) These resources have been created independently and with differing goals, but they are surprisingly compatible. They also differ primarily in the granularity of the semantic role labels. PropBank uses very generic labels such as Arg0 and Arg1, as in:

(B.1) *President Bush has approved duty-free treatment for imports of certain types of watches.*

REL: approved

Arg0: President Bush

Arg1: duty-free treatment for imports of certain types of watches.

NOTE The other numbered arguments in PropBank, Arg2 to Arg5, are quite verb-specific.

EngVallex uses non-numbered labels (e.g. ACT, PAT, ADDR, EFF and ORIG), which, with the exception of the first two, make them more descriptive, irrespective of the verb they are the argument of. In addition to providing several alternative syntactic frames and a set of semantic predicates, VerbNet marks the PropBank Arg0 as an Agent, and the Arg1 as a Theme. FrameNet labels them Grantor and Action, respectively, and puts them in the Grant Permission class. The additional semantic richness provided by VerbNet and FrameNet does not contradict PropBank and can be seen as complementary. Each of these resources is discussed in the succeeding subclauses, beginning with the most fine-grained one, FrameNet. The LIRICS project, Linguistic InFRastructure for Interoperable ResourCes and Systems, has made a serious study of these different frameworks and of the theoretical linguistics background. Their conclusions, including a set of Semantic Role definitions in the form of ISO data categories, are summarized in [Annex A](#) and can be found at

http://let.uvt.nl/general/people/bunt/docs/LIRICS_semrole.htm.

B.2 FrameNet

The FrameNet database is based on Fillmore's Frame Semantics,^{[18] [19] [21]} which asserts that much of the lexicon of any language is best understood as expressions that evoke both a state of affairs (a **semantic frame**) and the participants in it; these are given a frame-specific set of semantic role labels (**frame elements**). For example, the Apply_heat frame is evoked by the words such as *bake*, *barbecue*, *blanch*, *boil*, *braise*, *broil*, and *brown* (these are called **lexical units**); the frame elements of the Apply_heat frame include the Cook, the Food, and the Heating Instrument. More traditional labels (see [Table 1](#)) for the same roles might be Agent, Patient, and Instrument. The lexical units of Apply_heat all happen to be verbs, but a frame can have lexical units of any part of speech.¹⁾ For example, a more complex frame, Revenge, has lexical units such as *revenge.n*, *avenge.v*, *retaliate.v*, *vengeful.a*, and *vengeance.n*; the frame elements include the Offender, the Avenger, the Victim, the Injury, and the Punishment. A major goal of FrameNet is to document all the syntactic patterns in which each lexical unit can occur (its **valence**)

1) Many of the nouns in FrameNet denote events and, like *Achieving-first.invention.n*, *Assessing.evaluation.n*, and *Awareness.comprehension*, are derivationally related to verbs, but there are also other types of noun, such as *Make-agreement-on-action.treaty.n*, *Naturalfeatures.valley.n*, and *Part-inner-outer.exterior.n*.

by annotating example sentences from a corpus. For example, the following example sentence shows annotation with respect to the frame evoking expression *retaliated* in the Revenge frame; the frame elements Avenger, Injury, and Punishment are all annotated:

[Avenger The USA] RETALIATED [Injury against the harassment of its diplomats]

[Punishment by expelling 36 staff from the Iraqi embassy in Washington...]

FrameNet currently has 1 014 individual frames containing more than 11 500 lexical units, and of these, more than 6 500 have been annotated to document their valence.

A decision was made **not** to use any existing set of thematic roles for FrameNet because it was recognized that, for many frames (e.g. Revenge), none of the conventional thematic role names were adequate to represent the relations among the participants. Rather than force the roles of the frame onto the Procrustean bed of a small set of thematic roles, it was decided to define all the roles relative to each frame, and then add explicit frame-to-frame and frame element-to-frame element relations that would create a hierarchy, linking more specific frames to more general ones. The top-level frames of this hierarchy have frame elements that correspond to thematic roles (e.g. the frame *Intentionally_affect* has the frame elements Agent and Patient). The frame element *Apply_heat.Cook* inherits from *Intentionally_affect.Agent*, and *Apply_heat.Food* inherits from *Intentionally_affect.Patient*; this inheritance corresponds to strict subtyping. The advantage is that Cook can be given a much more specific definition within the *Apply_heat* frame and still be recognized as a type of Agent via the frame element relations, but at the cost of some complexity of representation. Sometimes there are simply no high-level roles that correspond to the more specific roles, and in these circumstances, the inheritance relations are simply not filled in. For example, the Similarity frame has frame elements Entity1 and Entity2, which do not correspond to any of the usual thematic roles (LIRICS would label these Pivot and Partner). The 1,0²⁾ frames are associated with well over 9 000 frame elements, most of which are linked to high-level 'thematic role' frame elements.^{[20]³⁾}

The Frame Elements for an individual Frame are classified into three levels, depending on how central they are to the definition of the frame:

- **core** [conceptually necessary for the definition of the frame (e.g. Cook and Food in *Apply_heat*), which, in the case of verbs, usually appear as arguments];
- **peripheral**, that is to say, not specific to the frame, but providing additional information, such as time and place, and similar to adjuncts;
- **extra-thematic**, that is to say, not part of the current frame, but related to another frame that frequently co-occurs with it. The database also contains relations such as 'requires' and 'excludes' between frame elements of the same frame, representing dependencies between them.

Unlike Levin's verb classes, lexical units are grouped into frames solely on the basis of having the same frame semantics, without regard to similarity of syntactic behaviour. Sets of verbs with similar syntactic behaviour can therefore appear in multiple frames, and a single FrameNet frame can contain sets of verbs with related senses but different subcategorization properties. FrameNet places a primary emphasis on providing rich, idiosyncratic descriptions of the semantic properties of lexical units in context, and on making explicit subtle differences in meaning. However, despite the different motivations, there are still many overlaps between verbs in the same Levin class and verbs associated with the same FrameNet frame. For instance, the Levin *Cooking* 45.3 class contains all of the FrameNet *Apply-heat* verbs, except for *singe*. It also includes a few additional, fairly infrequent verbs, many of which have to do with frying, such as *french-fry*, *oven-fry*, *oven-poach*, *overcook*, *overheat*, *pan-broil*, and *pan-fry*, as well as a few rare gems such as *parch*, *rissole*, *scallop*, and *schirr*. As might be expected, the greatest overlap between

2) Frame elements are given frame-specific names wherever possible, but there can be two distinct FEs in two different frames with the same name; in other words, FE names are only unique within frames. For example, where no suitably specific name suggests itself, FEs may be named Agent, Theme etc, even in lower-level frames.

3) Frame elements are given frame-specific names wherever possible, but there can be two distinct FEs in two different frames with the same name; in other words, FE names are only unique within frames. For example, where no suitably specific name suggests itself, FEs can be named Agent, Theme, etc., even in lower-level frames.

FrameNet and Levin classes occurs with the Levin classes that are most semantically coherent. It goes without saying that some Levin classes, such as *Braid 41.2.3: bob, braid, brush, clip, cold-cream, comb, condition, crimp, crop, and curl* are not intended to be semantically coherent, and they have little overlap with any FrameNet Frame. (See Reference [2] for a more detailed discussion of the similarities and differences between Levin's classes and FrameNet.)

B.3 VerbNet

B.3.1 Overview

VerbNet[13] [28] [29] [30] is midway between PropBank and FrameNet in terms of lexical specificity, and is closer to PropBank with regard to its close ties to syntactic structure. It consists of hierarchically arranged verb classes, inspired by and extended from Levin's classification of English verbs [Levin, 1993]. [74] The Levin classes have 240 classes, with 47 top-level classes and 193 second- and third-level classes. VerbNet has added almost 1 000 lemmas as well as over 200 more classes. There is now a fourth level of classes, and several additional classes at the other three levels. VerbNet adds to each Levin class an abstract representation of the syntactic frames, with explicit correspondences between syntactic positions and the semantic roles they express, as in *Agent REL Patient*, or *Patient REL into pieces for break*. (For other extensions of Levin, see also References [33] and [75]) The original Levin classes constitute the first few levels in the hierarchy, with each class subsequently refined to account for further semantic and syntactic differences within a class. In many cases, the additional information that VerbNet provides for each class has caused it to subdivide, or use intersections of, Levin classes. Each class and subclass is characterized extensionally by its set of verbs, and intensionally by a list of the arguments of those verbs and the syntactic and semantic information about them. The argument list consists of semantic roles, 24 in total, including Agent, Patient, Theme, and Experiencer (for the complete list, see Reference [66]), and possible selectional restrictions on the arguments that are expressed using binary predicates. In its current state, VerbNet makes use of the following:

- a) commonly used, coarse-grained roles like those of LIRICS (e.g. Agent);
- b) roles that are specific to certain classes of events, and which are intended to convey key semantic components of some verb classes (e.g. Topic, which is restricted to verbs of communication);
- c) roles that are in part syntactically motivated (e.g. Predicate, which is used for classes with predicative complements);
- d) roles that are distinguished by internal properties of the participant (e.g. [+animate]).

The semantic predicates describe the participants during various stages of the event expressed by the syntactic frame, and provide class-specific interpretations of the semantic roles. VerbNet now covers 3 965 verb lexemes with 471 classes. There are explicit links to similar entries in WordNet, OntoNotes groupings, FrameNet, and PropBank. A primary emphasis for VerbNet is the coherent syntactic and semantic characterization of the classes that will facilitate the acquisition of new class members based on observable syntactic and semantic behaviour.

B.3.2 Syntactic Frames

Each VerbNet class contains a set of syntactic descriptions, or syntactic frames, depicting the possible surface realizations of the argument structure. These include constructions such as transitive, intransitive, prepositional phrases, resultatives, and a large set of diathesis alternations listed by Levin as part of each verb class. Each syntactic frame is associated with an explicit thematic grid consisting of a set of semantic roles, such as Agent, Theme, and Location. The syntactic frame may specify other lexical items required for a particular construction or alternation. Semantic restrictions such as ANIMATE, HUMAN, and ORGANIZATION are used to constrain the types of semantic roles allowed in the classes. The 36 semantic types are taken originally from the EuroWordNet Interlingua, and can be viewed on the web. [67] They typically encompass literal meanings rather than metaphorical ones, and should be thought of as preferences rather than as hard constraints. Each syntactic frame may also be constrained in terms of which prepositions are allowed. Additionally, further restrictions may be imposed on semantic roles to indicate the syntactic nature of the constituent likely to be associated with

it. Levin classes are characterized primarily by Noun Phrase and Prepositional Phrase complements. Several additional classes based on work by Korhonen and Briscoe^[32] have been added to the original Levin classes, and many of these also include sentential complements. They refer only to the distinction between finite and nonfinite clauses, as in the various subclasses of Verbs of Communication.^[28]

B.3.3 Semantic predicates

Semantic predicates that denote the relations between participants and events are used to convey the key components of meaning for each class in VerbNet. The semantic information for verbs in VerbNet is expressed as a conjunction of semantic predicates such as MOTION, CONTACT, and CAUSE. As the classes may be distinguished by their temporal characteristics (e.g. Verbs of *Assuming a Position* vs. Verbs of *Spatial Configuration*), it is also necessary to convey information about when each of the predicates applies. In order to capture this information, semantic predicates are associated with an event variable, *e*, and often with START(*e*), END(*e*) or DURING(*e*) arguments, to indicate that the semantic predicate is in force either at the START, the END, or DURING the related time period for the entire event. Version 3.0 of VerbNet has 94 distinct semantic predicates, and an effort is currently under way to link the verb classes to the Omega ontology (0) and to create upper level nodes (0).

B.3.4 Verb class hierarchies

The classes of VerbNet are organized hierarchically and numbered in such a way that classes with shared class numbers have certain semantic and syntactic properties in common; for example, all classes numbered 9.1 to 9.10 are classes of verb that involve some kind of placement event, where one object is moved to a new location: *Put*, *Put_Spatial*, *Funnel*, *Put_Direction*, *Pour*, *Coil*, *Spray*, *Fill*, *Butter*, and *Pocket*. This hierarchical organization of the classes facilitates machine-learning for a variety of NLP applications. If training data are sparse for a particular class, it is possible to conflate this class with similar classes nearby in the hierarchy that share similar syntactic and semantic characteristics.

Each individual class is also hierarchical in the sense that classes can include one or more 'subclasses'. Each child subclass inherits all of the information from its parent class: this includes compatibility with the parent class's syntactic frames, thematic roles, and semantic and syntactic restrictions on those roles. The subclass also adds information in the form of additional syntactic frames, thematic roles and restrictions. Verb members of a subclass are therefore compatible with all the behavioural information listed in the parent class, and with the additional frames or thematic role information listed in the subclass. For example, in the *Put-9.1* parent class, the thematic roles of Agent, Theme, and Destination are listed along with the syntactic frames NP V NP PP.Destination (*I put the book on/under/near the table*) and NP V NP ADVP (*I put the book here/there*). All of the verbs listed in the parent class are compatible with these roles and syntactic frames. The child class, *Put-9.1-1*, lists an additional syntactic frame NP V NP (*I stashed the book*). The verbs listed in this subclass are not only compatible with all of the information listed in the parent class, but also with this additional syntactic frame. Conversely, the verbs listed in the parent class are not compatible with the frame in the child class. The hierarchical nature of each class allows for a refinement of the information represented about verbal behaviour beyond what was represented in Levin's original classification. In Levin's original classification, generalizations were made that applied to 'most verbs', whereas the addition of subclasses allows for the behaviour of each individual verb to be represented more precisely.

B.4 PropBank

B.4.1 Overview

By contrast with the objectives of FrameNet and VerbNet, the primary goal in developing the Proposition Bank, or PropBank, was not lexical resource creation, but the development of an annotated corpus to be used as training data for supervised machine-learning systems.^[39] The first PropBank release consists of 1 M words from the Wall Street Journal portion of the Penn Treebank II with predicate argument structures for verbs, using semantic role labels for each verb argument. Although the semantic role labels are purposely chosen to be quite generic and theory-neutral (Arg0, Arg1, etc.), they are still intended to consistently annotate the same semantic role across syntactic variations. It follows that the Arg1 or Patient in *John broke the window* is the same window that is annotated as the Arg1 in *The window*

broke, even though it is the syntactic subject in one sentence and the syntactic object in the other. The primary goal of PropBank is to supply consistent, simple, general-purpose labelling of semantic roles for a large quantity of coherent text to support the training of automatic semantic role labellers; the Penn Treebank has supported the training of statistical syntactic parsers in the same way. PropBank also provides a lexicon that lists, for each broad meaning of each annotated verb, its 'Frameset' [i.e. the possible arguments in the predicate and their labels (its 'roleset')] and all possible syntactic realizations. This lexical resource is used as a set of verb-specific guidelines by the annotators, and it is clearly quite similar in nature to FrameNet and VerbNet although at a more coarse-grained level.

B.4.2 Arguments/Adjuncts

The PropBank-numbered arguments are meant to be interpreted in a predicate-specific manner, whereas those in ArgM (Argument Modifier) have a global interpretation. PropBank is also more focused on literal meaning than FrameNet, and pays less attention to clearly marking metaphorical usages and support verb constructions.

Because of the difficulty of defining a universal set of semantic or thematic roles covering all types of predicates, PropBank defines semantic roles on a verb-by-verb basis. An individual verb's semantic arguments are numbered, beginning with 0. For a particular verb, Arg0 is generally the argument-exhibiting features of a prototypical Agent,^[14] while Arg1 is a prototypical Patient or Theme. No consistent generalizations can be made across verbs for the higher numbered arguments, although steps have been taken to consistently define roles across members of VerbNet classes. In addition to verb-specific numbered roles, PropBank defines several more general ArgM roles that can apply to any verb; they are similar to adjuncts, and include LOCation, EXTent, Adverbial, CAUse, TeMPoral, MaNneR, and DIRection.

A set of roles or roleset corresponds to a distinct usage of a verb. It is associated with a set of syntactic frames indicating allowable syntactic variations in the expression of that set of roles, the Frameset. A polysemous verb may have more than one Frameset, when the differences in meaning are distinct enough to require different sets of roles, one for each Frameset. The tagging guidelines include a verb-specific descriptor field for each role, such as *baker* for Arg0 and *creation* for Arg1 in the example below. These are intended for use during annotation and as documentation, and do not have any theoretical standing. In addition, each Frameset is complemented by a set of examples that attempt to cover the range of syntactic alternations afforded by that usage. The example below the collection of Frameset entries for a verb is referred to as a *Frame File*. Furthermore, the neutral, generic labels facilitate mapping between PropBank and other more fine-grained resources such as VerbNet and FrameNet, and also Lexical-Conceptual Structure and Prague Tectogramatics.^[46] While most rolesets have between two and four numbered roles, as many as six can appear, in particular for certain verbs of motion. For more details, see Reference ^[39] and the online Frame Files.^[67]

(B.2) *Today whole grains are freshly ground every day and baked into bread.*

Table B.1 — Roleset bake.01

ROLES for <i>bake.01</i> Roleset	ANNOTATION: REL: <i>bake.01</i>
Arg0:baker	
Arg1:creation	Arg1 into bread
Arg2:source	Arg2: whole grains
Arg3:benefactive	

B.4.3 The argument/adjunct distinction in PropBank

The use of numbered arguments highlights PropBank's approach to the argument/adjunct distinction; for example, it uses unique predicate-specific numerical IDs to represent arguments, but a fixed set of global (non-predicate-specific) labels to represent adjuncts. As illustrated in (2), the arguments (Arg0, Arg1, Arg2) are represented by a numerical ID while the adjunct (ArgM-TEMP) is represented by a global label from a fixed pool of labels.

(B.3) [Arg1 The projector] was [REL broken] by [Arg0 Larry] [Arg2 with a hammer] [ArgM-TEMP yesterday].

The use of numerical IDs for arguments can only be interpreted properly because the arguments are individualised to the predicates, and because they are unique for a given predicate. It does not make sense for adjuncts because they are repeatable, and are not individualised. A consequence of this is that these unique IDs need to be interpreted differently for each different predicate, and even for different senses of a predicate. This is illustrated in [Table B.2](#):

Table B.2 — Predicate-specific numerical IDs for arguments

	Arg0	Arg1	Arg2	Arg3	Arg4	Arg5
loiter.02	loiterer					
leak.01	thing leaking	substance leaked				
replace.01	replacer	old thing	New thing			
translate.01	translator	thing changing	end stage	start state		
increase.01	causer of increase	thing increasing	Amount of increase	start point	end point	
send.03	causer of motion	entity in motion	extent of motion	start point	end point	Direction

However, adjuncts are represented by a set of global labels that are equally applicable to all predicates. Note that, strictly speaking, STR, NEG, MOD, and DIS are not adjuncts, and are included for more pragmatic reasons.

Table B.3 — Global labels for adjunct

Role	Description	Role	Description
ADV	Adverbial	LOC	Locatives
CAU	Cause clauses	NEG	Negation
DIR	Directionals	PNC	Purpose, not reason
DIS	Discourse markers	PRD	Secondary predicate
EXT	Extent markers	REC	Reciprocals
LOC	Locatives	STR	Stranded
MNR	MNR markers	TMP	Temporal markers

FrameNet has a different strategy for encoding the argument/adjunct distinction. FrameNet considers all arguments for a predicate as elements in a semantic frame. Arguments and adjuncts are not encoded differently, as they are in the PropBank, but FrameNet does make a distinction between *core* and *non-core* arguments; this roughly parallels the argument/adjunction distinction. Core and non-core arguments are explicitly listed in each of the semantic frames defined for FrameNet; for example, in the Commerce_Buy frame, the core arguments are Buyer and Goods, and the non-core arguments are Manner, Means, Money, Place, Purpose, Rate, Reason, Recipient, Seller, Time, and Unit.

In spite of its success in facilitating the training of semantic role labelling (SRL), there are several ways in which PropBank could be more effective, as discussed below. PropBank lacks much of the information that is contained in VerbNet, including information about selectional restrictions, verb semantics, and inter-verb relationships, and we have therefore created one mapping between VerbNet and PropBank, and another between VerbNet and FrameNet; this will allow us to use the machine-learning techniques that have been developed for PropBank annotations to generate more semantically abstract VerbNet and FrameNet representations. This, too, is discussed below.

B.4.4 Limitations to a predicate-specific approach

There is a widely acknowledged lack of consensus in the community on a specific set of semantic role labels. PropBank avoids this issue by using theory-agnostic labels (Arg0, Arg1, ..., Arg5) and by defining

those labels as having verb-specific meanings. Under this scheme, PropBank can avoid making any claims about how any one verb's arguments relate to other verbs' arguments or about general distinctions between verb arguments and adjuncts. However, there are several limitations to this approach. The first is that it can be difficult to make role label-based inferences and generalizations based on role models that are only meaningful with respect to a single verb. Since each role label is verb-specific, we cannot confidently determine when two verbs' arguments have the same role; and since no encoded meaning is associated with each tag, we cannot make generalizations across verb classes. By contrast, the use of a shared set of role labels, as in VerbNet, facilitates both inferencing and generalization. An additional issue with PropBank's verb-specific approach is that it can make training automatic semantic role labelling (SRL) systems more difficult. As with FrameNet, a vast amount of data would be needed to train the verb-specific (or frame-specific) models that are theoretically mandated by the fine-grained role labels. For the most part, researchers who use PropBank as training data ignore the 'verb-specific' nature of the labels, and instead build a single model for each numbered argument (Arg0, Arg1, ..., Arg5). Given the correspondence between Arg0/Arg1 and Dowty's Proto-Agent/Proto-Patient, and the fact that they correspond to 85 % of the arguments, it is not surprising that this is effective. The ArgM's are also labelled quite consistently, but arguments Arg2 to Arg5 are very overloaded, and automatic Semantic Role Labelling performance drops significantly on them. A final limitation arises from the genre-specific nature of the training corpus, which initially consisted entirely of Wall Street Journal articles. This has since been expanded with DARPA-GALE funding to include Broadcast News, Broadcast Conversation, Newsgroups, and Weblogs, but significant additional quantities of corpora will still be needed to train a truly robust system. This issue is reflected in the relatively poor performance of most state-of-the-art SRL systems when tested on a novel genre, the Brown corpus, during CoNLL 2005. For example, the SRL system, described in References [43] and [44], achieves an F-score of 81 % when tested on the same genre as it is trained on (WSJ); but that score falls to 68,5 % when the same system is tested on a different genre (the Brown corpus). Better techniques for generalizing the semantic role-labelling task are still needed in addition to the new DARPA-GALE genres, as in Reference [54]. It would also be advantageous to be able to merge the FrameNet- and PropBank-labelled instances to create a much larger, more diverse, and yet still coherent, training corpus.

B.5 EngVallex

In the same way that PropBank was developed with the intention of mapping to the Penn Treebank II phrase structure parses, EngVallex^[16] was created to complement the Prague Dependency-style annotation of English, namely the texts of the WSJ portion of the Penn Treebank. It is based on the theory of valency found in the Functional Generative Description dependency framework,^[49] and the main objective was to ensure consistency of semantic role annotation in the corpus. The roles, labels, lexicon structure and annotation have been modelled on the PDT-Vallex, a valency frame lexicon for Czech.^[52] The Functional Generative Description Valency Theory (FGDVT) is dependency-oriented: it 'operates' on the tectogrammatical layer, and combines the syntactic and semantic approach for distinguishing valency elements. The verb is considered to be the core of the sentence. Within this approach, syntactic as well as semantic criteria are used to identify verbal complementations. It is assumed that every (semantic) verb (and potentially noun, adjective, and adverb) has certain subcategorization requirements that can be specified in the valency frame of the given unit. The verbal valency modifications are specified along two axes. The first axis concerns the (general) opposition between inner participants (arguments) and free modifications (adjuncts); this distinction is based on certain operational criteria relating to

- a) the possibility of the appearance of the same type of complementation with the same verb, and
- b) the possibility of the occurrence of the given complementation (principally) with all verbs.

The other axis relates to the distinction between (semantically) obligatory and optional complementations of the given unit; the determination of semantic (non)-obligatoriness is again based on a certain operational criterion (the so-called dialogue test) proposed by J. Panevová. Five inner participants (arguments) are distinguished: Actor/Bearer (ACT), Patient (PAT), Addressee (ADDR), Origin (ORIG), and Effect (EFF). The repertory of free modifications (adjuncts) is much larger than that of arguments; for example, FGD distinguishes about 50 types of adjuncts. Adjuncts are always determined semantically: their set might be divided into several subclasses, such as temporal (TWHEN, TSIN, TTILL, TFL, TFHL, THO, TPAR, TFRWH, and TOWH), local (LOC, DIR1, DIR2, and DIR3), causal (e.g. CAUS for Cause, AIM,

and CRIT for ‘according to’) and other free modifications (e.g. MANN for general ‘Manner’, ACMP for Accompaniment, EXT for Extent, MEANS, INTF for Intensifier, and BEN for Benefactor). Only arguments (obligatory or optional) and obligatory adjuncts are considered to be an integral part of any verbal valency frame.

The Prague Czech-English Dependency Treebank (PCEDT^[25]) belongs to the family of Prague treebanks (PDT, PADT, PEDT). The PCEDT is a parallel Czech-English treebank with valency information annotated in both languages. Both this treebank and the PDT are annotated on three layers. On the lowest, morphological layer, the lexical entry (usually represented by a lemma), and values of morphological categories (e.g. person, number, tense, gender, voice, and aspect) are assigned to each word. At the analytical layer, a sentence is represented as a dependency tree. Nodes of the tree represent tokens (i.e. word forms and punctuation marks) as they are found in the original sentence; no nodes are added or deleted. Where appropriate, edges usually represent a relation of formal dependency where it makes sense to do so. An analytical function capturing the type of dependency relation between the child and its parent is also added. The highest (or ‘semantically deepest’) layer, which is called the ‘tectogrammatical layer’, captures the deep (underlying) structure of a sentence. Nodes represent only autosemantic words, but synsemantic (i.e. auxiliary) words and punctuation marks are not represented by nodes as they only affect values of attributes of the autosemantic words they are attached to. At this layer, several attributes are assigned to each node, one of the most important being the (deep) functor that captures the tectogrammatical function of a dependent relative to its governor. It is here that the valency comes into play, with each verb (the governor) and its dependents annotated for verb sense and a valency frame assigned to each occurrence of every verb in the corpus. The extracted valency frames also form a valency lexicon: PDT-Vallex for Czech and EngVallex for English. Since the data annotated consists of the Penn Treebank WSJ texts, each occurrence of a verb is also linked to its PropBank annotation and appropriate role in the frame file.

[Table 2](#) provides an overview of the current status of the four afore-mentioned resources that have extensive annotated data.

B.6 Linking different frameworks

B.6.1 Overview

With the dual goal of being able both to merge PropBank and FrameNet training data and to map back and forth between PropBank, VerbNet, and FrameNet labellings for annotated instances, type-to-type mappings have been made between PropBank and VerbNet and between VerbNet and FrameNet. These mappings have been used to leverage a mapping of the PropBank-annotated instances to the relevant VerbNet classes and semantic role labels. Steps are being taken to extend this instance mapping to FrameNet. This project is called Semlink.

B.6.2 VerbNet-PropBank

The mapping between VerbNet and PropBank is in two parts: a lexical mapping and an annotated corpus. The lexical mapping is responsible for specifying the potential mappings between PropBank and VerbNet for a given word; but it does not specify which of these mappings should be used for any given occurrence of the word. That is the job of the annotated corpus, which, for any given instance, gives the specific VerbNet mapping and semantic role labels. This can be thought of as a form of sense-tagging. Where a PropBank frame maps to several VerbNet classes, they can be thought of as more fine-grained senses, and labelling with the class label corresponds to providing a sense tag label. The lexical mapping was used to automatically predict VerbNet classes and role labels for each instance. Where the resulting mapping was one-to-many, the correct mapping was selected manually (0). The usefulness of this mapping for improving SRL on new genres is discussed in Reference [\[40\]](#).

B.6.3 VerbNet-FrameNet

The SemLink VerbNet/FrameNet mapping is in three parts. The first part is a many-to-many mapping of VerbNet Classes and FrameNet Frames. It is many-to-many in that a given FrameNet lexical unit can map to more than one VerbNet member, and a given VerbNet member can more often map to more than