

VerbNet Representations: Subevent Semantics for Transfer Verbs

Susan Windisch Brown

University of Colorado
brownsw@colorado.edu

Julia Bonn

University of Colorado
bonnj@colorado.edu

James Gung

University of Colorado
james.gung@colorado.edu

Annie Zaenen

Stanford University
azaenen@stanford.edu

James Pustejovsky

Brandeis University
jamesp@cs.brandeis.edu

Martha Palmer

University of Colorado
umpalmer@colorado.edu

Abstract

This paper announces the release of a new version of the English lexical resource VerbNet with substantially revised semantic representations designed to facilitate computer planning and reasoning based on human language. We use the transfer of possession and transfer of information event representations to illustrate both the general framework of the representations and the types of nuances the new representations can capture. These representations use a Generative Lexicon-inspired subevent structure to track attributes of event participants across time, highlighting oppositions and temporal and causal relations among the subevents.

1 Introduction

Many natural language processing tasks have seen rapid advancement in recent years using deep learning methods; however, those tasks that require precise tracking of event sequences and participants across a discourse still perform better using explicit representations of the meanings of each sentence or utterance. To be most useful for automatic language understanding and generation, such representations need to be both automatically derivable from text and reasonably formatted for computer analysis and planning systems. For applications like robotics or interactions with avatars, commonsense inferences needed to understand human language directions or interactions are often not derivable directly from the utterance. Tracking intrinsic and extrinsic states of entities, such as their existence, location or functionality, currently requires explicit statements with precise temporal sequencing.

In this paper, we describe new semantic representations for the lexical resource VerbNet that provide this sort of information for thousands of

verb senses and introduce a means for automatically translating text to these representations. We explore the format of these representations and the types of information they track by thoroughly examining the representations for transfer of possessions and information. These event types are excellent examples of complex events with multiple participants and relations between them that change across the time frame of the event. By aligning our new representations more closely with the dynamic event structure encapsulated by the Generative Lexicon, we can provide a more precise subevent structure that makes the changes over time explicit (Pustejovsky, 1995; Pustejovsky et al., 2016). *Who has what when* and *who knows what when* are exactly the sorts of things that we want to extract from text, but this extraction is difficult without explicit, computationally-tractable representations. These event types also make up a substantial portion of VerbNet: 37 classes of verbs deal with change of possession and transfer of information out of VerbNet’s 300+ classes, covering 810 verbs.

2 Background

The language resource VerbNet (Kipper et al., 2006) is a hierarchical, wide-coverage verb lexicon that groups verbs into classes based on similarities in their syntactic and semantic behavior (Schuler, 2005). Each class in VerbNet includes a set of member verbs, the thematic roles used in the predicate-argument structure of these members (Bonial et al., 2011), and the class-specific selectional preferences for those roles. The class also provides a set of typical syntactic patterns and corresponding semantic representations. A verb can be a member of multiple classes; for example, *run* is a member of 8 VerbNet classes, including the run-51.3.2 class (*he ran to the store*)

and the function-105.2.1 class (*the car isn't running*). These memberships usually correspond to coarse-grained senses of the verb. The resource was originally based on Levin's (1993) analysis of English verbs but has since been expanded to include dozens of additional classes and hundreds of additional verbs and verb senses.

VerbNet representations previously formed the basis for Parameterized Action Representation (PAR) providing a conceptual representation of different types of actions (Badler et al., 1999). These actions involve changes of state, changes of location, and exertion of force and can be used to animate human avatars in a virtual 3D environment (R. Bindiganavale and Palmer, 2000). They are particularly well suited for motion and contact verb classes, providing an abstract, language-independent representation (Kipper and Palmer, 2000). The more precise temporal sequencing described here is even more suitable as a foundation for natural language instructions and human-robot or human-avatar interactions.

2.1 VerbNet

VerbNet has long been used in NLP for semantic role labeling and other inference-enabling tasks (Shi and Mihalcea, 2005; Giuglea and Moschitti, 2006; Loper et al., 2007; Bos, 2008). In addition, automatic disambiguation of a verb's VerbNet class has been used as a stand-in for verb sense disambiguation (Abend et al., 2008; Brown et al., 2014; Croce et al., 2012; Kawahara and Palmer, 2014).

VerbNet's semantic representations use a Davidsonian first-order-logic formulation that incorporates the thematic roles of the class. Each frame in a class is labeled with a flat syntactic pattern (e.g., NP V NP). The "syntax" that follows shows how the thematic roles for that class appear in that pattern (e.g., Agent V Patient), much like the argument role constructions of Goldberg (2006). A previous revision of the VerbNet semantic representations made the correspondence of these patterns to constructions more explicit by using a common predicate (i.e., `path_rel`) for all caused-motion construction frames (Hwang, 2014). At the request of some users, we are substituting more specific predicates for the general `path_rel` predicate, such as `has_location`, `has_state` and `has_possession`, although the subevent patterns continue to show the commonality across

these caused-motion frames.

Each frame also includes a semantic representation that uses basic predicates to show the relationships between the thematic role arguments and to track any changes over the time course of the event. Thematic roles that appear in the "syntax" should always appear somewhere in the semantic representation. Overall, this linking in each frame of the syntactic pattern to a semantic representation is a unique feature of VerbNet that emphasizes the close interplay of syntax and semantics.

2.2 Revision of the Semantic Representations

VerbNet's old representations included an event variable E as an argument to the predicates. Representations of states were indicated with either a bare E, as for the `own-100` class: `has_possession(E, Pivot, Theme)`, or `During(E)`, as for the `contiguous_location-47.8` class (*Italy borders France*): `contact(During(E), Theme, Co-Theme)`. Most classes having to do with change, such as changes in location, changes in state and changes in possession, used a `path_rel` predicate in combination with `Start(E)`, `During(E)`, and `End/Result(E)` to show the transition from one location or state to another (1).

- (1) *The rabbit hopped across the lawn.*
 Theme V Trajectory
`motion(during(E), Theme)`
`path_rel(start(E), Theme, ?Initial location1, CH_OF_LOC, prep)`
`path_rel(during(E), Theme, Trajectory, CH_OF_LOC, prep)`
`path_rel(end(E), Theme, ?Destination, CH_OF_LOC, prep)`

Efforts to use VerbNet's semantic representations (Zaenen et al., 2008; Narayan-Chen et al., 2017), however, indicated a need for greater consistency and expressiveness. We have addressed consistency on several fronts. First, all necessary participants are accounted for in the representations, whether they are instantiated in the syntax, incorporated in the verb itself (e.g., *to drill*), or simply logically necessary (e.g., all entities that change location begin in an initial location, whether it is commonly mentioned or not).

¹A question mark in front of a thematic role indicates a role that appears in the syntax in some frames for the class but not in this particular frame.

Second, similar event types are represented with a similar format; for example, all states are represented with E, never with *During*(E). Finally, predicates are given formal definitions that apply across classes.

In order to clarify what is happening at each stage of an event, we turned to the Generative Lexicon (Pustejovsky, 1995) for an explicit theory of subevent structure. Classic GL characterizes the different Aktionsarten in terms of structured subevents, with states represented with a simple *e*, processes as a sequence of states characterizing values of some attribute, $e_1 \dots e_n$, and transitions describing the opposition inherent in achievements and accomplishments. In subsequent work within GL, event structure has been integrated with dynamic semantic models in order to more explicitly represent the attribute modified in the course of the event (the location of the moving entity, the extent of a created or destroyed entity, etc.) as a sequence of states related to time points or intervals. This Dynamic Event Model (Pustejovsky and Moszkowicz, 2011; Pustejovsky, 2013) explicitly labels the transitions that move an event from frame to frame.

Applying the Dynamic Event Model to VerbNet semantic representations allowed us refine the event sequences by expanding the previous tripartite division of *Start*(E), *During*(E), and *End*(E) to an indefinite number of subevents. These numbered subevents allow very precise tracking of participants across time and a nuanced representation of causation and action sequencing within a single event. In the general case, e_1 occurs before e_2 , which occurs before e_3 , and so on. We've introduced predicates that indicate temporal and causal relations between the subevents, such as **cause**(e_i , e_j) and **co-temporal**(e_i , e_j).

We have made other refinements suggested by the GL Dynamic Event Model. For example, we greatly expanded the use of negated predicates to make explicit the opposition occurring in events involving change: e.g., *John died* is analyzed as the opposition $\langle \mathbf{alive}(e_1, \text{Patient}), \neg \mathbf{alive}(e_2, \text{Patient}) \rangle$. Compare the new representation for changes of location in (2) to (1) above. In (2), we use the opposition between **has_location** and $\neg \mathbf{has_location}$ to make clear that once the Theme is motion (in e_2), it is no longer in the *Initial_location*. In order to distinguish the event type associated with a semantic predicate, we in-

troduced a new event variable, \ddot{e} , to distinguish a process from other types of subevents, such as states. For example, see the motion predicate in (2).

- (2) *The rabbit hopped across the lawn.*
 Theme V Trajectory
has_location(e_1 , Theme, ?Initial_Location)
motion(\ddot{e}_2 , Theme, Trajectory)
 $\neg \mathbf{has_location}$ (e_2 , Theme, ?Initial_location)
has_location(e_3 , Theme, ?Destination)

Although people infer that an entity is no longer at its initial location once motion has begun, computers need explicit mention of this fact to accurately track the location of an entity. Similarly, some states hold throughout an event, while others do not. Our new representations make these distinctions clear, where *pre-event*, *while-event*, and *post-event* conditions are distinguished formally in the representation.

Elsewhere (Brown et al., 2018), we discuss in more detail the Dynamic Event Model, show the effect of the new subevent structure on the interpretation of the role of the Agent, and give further examples of the new change of location and change of state representations.

3 Change of Possession

In this section, we closely examine the representations for events involving changes in possession. These representations illustrate the greater clarity and flexibility we have gained by adopting the conventions described in section 2. They also show some of the choices we have made to capture the underlying semantics while maintaining a connection to the varying surface forms. We discuss both one-way transfers (*give*) and two-way transfers (*sell*). We also address the different perspectives verbs can impose on a transfer event, such as the difference between *Mary gave John the book* and *John obtained the book from Mary*, in which the Agent of the event is the Source or the Recipient of the item, respectively. These variations have interesting analogs in the Transfer of Information classes (Fig. 1), which we discuss in Section 4.

The semantic representations for changes of possession in VerbNet assume a literal, non-metaphoric use of the verbs in question. Metaphor may select only some of the source domain's participants or entailments. For example, *She stole*

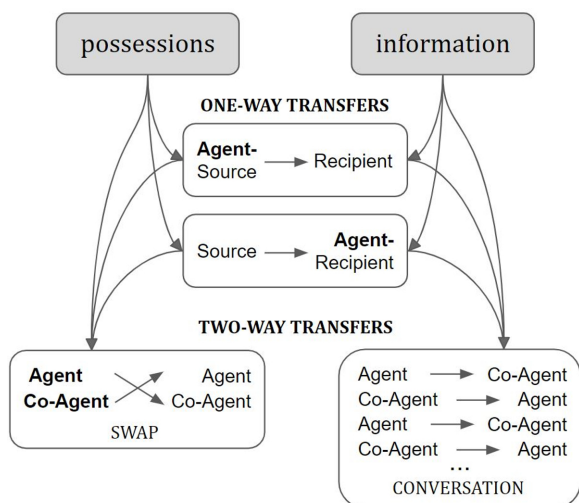


Figure 1: Primary distinctions made in the VerbNet representations for events involving transfer

John's car entails that John no longer has possession of his car, whereas *She stole John's heart* does not entail his loss of a vital organ. An analysis of VerbNet classes in terms of their application to figurative language (Brown and Palmer, 2012) showed that some classes concern only metaphoric uses of their member verbs (e.g., *calibratable_cos-45.6.1*), with semantic representations that directly represent the figurative meaning without reference to the source domain. Many classes, however, were shown to refer to literal uses of the verbs, although it was suggested that transformations or re-interpretations of the semantic representations could be possible.

3.1 Previous Representations

The previous model allowed only three temporal subevent periods: Start, During, and End. For both Change of Possession and Transfer of Information classes, each possession received one **path_rel** for the Start period and one for the End period, allowing one clear owner per period. For Change of Possession, it was reasonable to assume that possession transferred fully during the event, and as such, information about who did not possess a thing at any point could have been inferred through a rule. This model was sufficient for Change of Possession classes in and of themselves, but failed to capture any contrast with Transfer of Information classes, for which this assumption does not hold.

The **cause** predicate included arguments for an Agent or Causer (no other thematic roles were allowed), and the overall event E. This was sufficient

for one-way transfers in which one party was responsible for initiating the entire change, but was insufficient when more than one transfer occurred. There was no way to show that one party could initiate one transfer while another party initiated another. Two-way transfer representations either attributed all causation to one party, or omitted the **cause** predicate entirely. The ability to omit the predicate led one to wonder why it was ever necessary to include it.

3.2 New Representations

Three predicates form the core of the change of possession representations:

- **has_possession**(e, [slot-1], [slot-2])
- **transfer**(e, [slot-1], [slot-2], [slot-3])
- **cause**(e_i, e_j)

We define **has_possession** broadly as involving ownership or control over a thing; e.g., *I have a pencil* can mean either you own a pencil or you (possibly temporarily) have use of a pencil. Within the predicate, slot-1 is reserved for the possessor and can take thematic roles Source, Recipient, Goal, Agent, and Co-Agent. Slot-2 is reserved for the possession, and can take roles Theme and Asset.

Transfer is now a causative predicate, describing an event in which possession of a thing transfers from one possessor to another. All three participants are given as arguments. Slot-1 is reserved for the possessor who initiates the transfer, and can take thematic roles Agent and Co-Agent. Slot-2 is reserved for the possession (Theme or Asset), and slot-3 is reserved for the other possessor (Source, Recipient/Goal, Agent, and Co-Agent).

The order of arguments within this predicate often aligns with the temporal order of possession, but this is incidental. Sometimes, an Agent who is initiating a **transfer** is the recipient of that transfer; in these cases, the Agent will still occupy slot-1, even though they end up with possession last. It is also possible for an Agent to occupy slot-3 if another party (Co-Agent) is initiating the **transfer**. The subevent numbering of the **has_possession** predicates before and after the **transfer** provide a full description of the temporal order of possession.

The new basic representation is shown in (3).

- (3) **has_possession**(e_1 , Source, Theme)
–has_possession(e_1 , Recipient, Theme)

transfer(e_2 , Source, Theme, Recipient)
has_possession(e_3 , Recipient, Theme)
¬has_possession(e_3 , Source, Theme)
cause(e_2 , e_3)

This representation contains an initial state subevent, a **transfer** subevent, and a resulting state subevent. **Cause**(e_2 , e_3) tells us that the **transfer** triggers the resulting state. The opposing **¬has_possession** predicates show without a doubt that the Source stops having possession as soon as the transfer occurs, and the Recipient does not take possession until then. This allows for clear automatic tracking of an entity's ownership status and provides an important contrast with the new Transfer of Information representations. It will also allow coverage of cases of shared ownership of possessions, if VerbNet expands in that direction.

3.3 Change of Possession Variations

Agents as Sources or Recipients: Depending on the class, an Agent may function as Source or Recipient. In the old representations, some classes ended up including as core roles both an Agent and a Recipient, or an Agent and a Source, even if those roles always overlapped in the syntax. This was likely due to pressure to include in the class thematic roles that were projected by the main predicates, **path_rel**, **cause** and **transfer**. In the new model, we let Agent stand in for whichever role it overlaps throughout the representation. This eliminates the need for the **equals** predicate, and has allowed us to eliminate syntactically redundant roles from the class role inventories.

Six classes demonstrate one-way transfers in which the entity who starts with possession initiates giving that possession away: *cheat*-10.6.1, *contribute*-13.2, *equip*-13.4.2, *fulfilling*-13.4.1, *future_having*-13.3, and *give*-13.1-1. In example (4) from *fulfilling*-13.4.1, Agent replaces Source throughout.

(4) *Brown presented a plaque to Jones*
 Agent V Theme Recipient
has_possession(e_1 , Agent, Theme)
¬has_possession(e_1 , Recipient, Theme)
transfer(e_2 , Agent, Theme, Recipient)
has_possession(e_3 , Recipient, Theme)
¬has_possession(e_3 , Agent, Theme)
cause(e_2 , e_3)

Five classes demonstrate one-way transfers in which an entity who does not have possession of

a thing initiates taking that thing from the original possessor: *berry*-13.7, *deprive*-10.6.2, *obtain*-13.5.2, *rob*-10.6.4, and *steal*-10.5. The example from *steal*-10.5 in (5) shows how Agent replaces Recipient.

(5) *They stole the painting from the museum*
 Agent V Theme Source
has_possession(e_1 , Source, Theme)
¬has_possession(e_1 , Agent, Theme)
transfer(e_2 , Agent, Theme, Source)
has_possession(e_3 , Agent, Theme)
¬has_possession(e_3 , Source, Theme)
cause(e_2 , e_3)

Four main classes and two additional subclasses belonging to classes listed above demonstrate two-way transfers: *exchange*-13.6.1, *get*-13.5.1, *invest*-13.5.4, and *pay*-68, as well as *give*-13.1-1 and *obtain*-13.5.2-1. In the following example from *exchange*-13.6.1, note the new handling of subevents, **cause**, and the argument structure of **transfer**. In e_2 , the Agent initiates the transfer of the Theme, and in e_3 , the Co-Agent initiates the transfer of the Co-Theme. Subevent e_2 causes the resulting possession states of the Theme, and e_3 causes the resulting possession states of the Co-Theme.

(6) *Gwen exchanged the dress for a shirt*
 Agent V Theme Co-Theme
has_possession(e_1 , Agent, Theme)
¬has_possession(e_1 , ?Co-Agent, Theme)
has_possession(e_1 , ?Co-Agent, Co-Theme)
¬has_possession(e_1 , Agent, Co-Theme)
transfer(e_2 , Agent, Theme, ?Co-Agent)
transfer(e_3 , ?Co-Agent, Co-Theme, Agent)
has_possession(e_4 , ?Co-Agent, Theme)
¬has_possession(e_4 , Agent, Theme)
has_possession(e_5 , Agent, Co-Theme)
¬has_possession(e_5 , ?Co-Agent, Co-Theme)
cause(e_2 , e_4)
cause(e_3 , e_5)

Substitute-13.6.2 used to be included in this group, but since it was specifically split off from *exchange*-13.6.1 to deal with a two-way exchange of location (i.e., two entities change places with each other), we are now treating it purely as a Change of Location class rather than Change of Possession. When compared with (6), example (7) from *substitute*-13.6.2 highlights the distinctions

we are able to achieve using the new Change of Location vs. Change of Possession treatments.

- (7) *One bell ringer swapped places with another*
 Theme V Location Co-Theme
has_location(e_1 , Theme, Location_I)
has_location(e_2 , Co-Theme, Location_J)
motion(\ddot{e}_3 , Theme, Trajectory)
–has_location(e_3 , Theme, Location_I)
motion(\ddot{e}_4 , Co-Theme, Trajectory)
–has_location(e_4 , Co-Theme, Location_J)
has_location(e_5 , Theme, Location_J)
has_location(e_6 , Co-Theme, Location_I)
cause(\ddot{e}_3 , e_5)
cause(\ddot{e}_4 , e_6)

Additional predicates: Several subgroups within Change of Possession use additional predicates to depict additional semantics. Future-having-13.3 and berry-13.7 both take an **irrealis**(e) predicate to show that the **transfer** and resulting states are intended, but not guaranteed to have taken place yet. **Irrealis**'s single argument is a subevent number, and one predicate is given per qualifying subevent. Another additional predicate is used in the get-13.5.1, give-13.1-1, obtain-13.5.2, and pay-68 classes. These all involve two-way transfers in which a Theme is exchanged for an Asset, where the Asset is the cost of the Theme, represented as **cost**(Theme, Asset). Finally, rob-10.6.4 and steal-10.5 both involve an Agent/Recipient who initiates taking a possession in an illegal manner. The representations include a **manner**(e , Illegal, Agent) predicate which, for this usage, takes Illegal as a constant.

4 Transfer of Information

4.1 Previous Representations

In the old model, the only consistent difference between Transfer of Information and Change of Possession in terms of predicates and representation structure lay within **path_rel**, which contained a constant called either TR_OF_INFO or CH_OF_POSS, respectively. Like Change of Possession, only one **path_rel** was provided per temporal period, allowing only one clear possessor per period. Unfortunately, this failed to capture the important distinctions that knowledge is generally not lost when communicated, and one party's possession and communication of knowledge is no

guarantee that another party doesn't already possess it too.

4.2 New Representations

Two new predicates describe Transfer of Information:

- **has_information**(e , [slot-1], [slot-2])
- **transfer_info**(e , [slot-1], [slot-2], [slot-3])

These mirror the predicates used in Change of Possession in terms of their argument slots and functions, excepting that slot-2 may take Theme or Topic but not Asset. Topic is used most commonly for verbal information, while Theme is reserved for non-verbal information, which often reflects assent or emotional states.

The basic representation in (8) differs from Change of Possession in terms of the boundaries on possession before and after the **transfer_info** subevent. Here, by leaving the Recipient's possession status underspecified in e_1 , we make no claims about whether or not the Recipient already knew the information at the beginning of the event. By marking the Source's possession status with a big E, we assert that the Source maintains possession of the information throughout the event, even after the **transfer_info** communication subevent.

- (8) **has_information**(E, Source, Topic)
transfer_info(e_1 , Source, Topic, Recipient)
has_information(e_2 , Recipient, Topic)
cause(e_1 , e_2)

4.3 Transfer of Information Variations

One-way transfers: Just as with Change of Possession, Transfer of Information classes may involve an Agentive Source or Agentive Recipient. The basic representations for these types alternate from the basic Transfer of Information representation in the same way demonstrated above, with Agent replacing either Source or Recipient throughout. The vast majority of Transfer of Information classes are of the Agentive Source type, including advise-37.9, complain-37.8, confess-37.10, crane-40.3.2, curtsy-40.3.3, initiate_communication-37.4.2, inquire-37.1.2, instr_communication-37.4.1, interrogate-37.1.3, lecture-37.1.1, manner_speaking-37.3, nonverbal_expression-40.2, overstate-37.12, promise-37.13, say-37.7, tell-37.2, transfer_msg-37.1.1, and wink-40.3.1. Just one class, learn-14, features an Agentive Recipient.

Two-way transfers: The two-way Transfer of Information classes, *chit_chat-37.6* and *talk-37.5*, differ from the two-way Change of Possession classes in several ways. Most notably, they are not limited to a single transfer in each direction; instead, a sequence of transfers repeats back and forth between the two participants an unspecified number of times. The subevent ordering is changed so that the state resulting from one **transfer_info** occurs before the next **transfer_info** begins. The repeated turn-taking is expressed using the **repeated_sequence** predicate, which may take as many subevent arguments as necessary to capture the full span of the repeated behavior. The example in (9) is from *chit_chat-37.6*.

(9) *Susan chitchatted with Rachel about the problem*

Agent V Co-Agent Topic
has_information(E, Agent, Topic_I)
has_information(E, Co-Agent, Topic_J)
transfer_info(e_1 , Agent, Topic_I, Co-Agent)
has_information(e_2 , Co-Agent, Topic_I)
transfer_info(e_3 , Co-Agent, Topic_J, Agent)
has_information(e_4 , Agent, Topic_J)
cause(e_1, e_2)
cause(e_3, e_4)
repeated_sequence(e_1, e_2, e_3, e_4)

Additional predicates and selectional restrictions: Several subgroups within Transfer of Information capture further semantic details using either additional predicates or specialized selectional restrictions on class roles. Two classes feature verbs of asking: *inquire-37.1.2* and *interrogate-37.1.3*. These classes take a Topic role with a selectional restriction [+question], which helps clarify that the communication event taking place regards the question and never the response. *Manner_speaking-37.3* and *nonverbal_expression-40.2* both feature verbs that describe the manner of communication. The representations use another **manner** predicate, this time with a verb-specific role V_Manner in place of a constant. *Instr_communication-37.4* features verbs that describe an instrument used to communicate (e.g., *phone*), and uses **utilize**(e , Agent, V_Instrument) to convey this.

Two subgroups use Theme with selectional restriction [+nonverbal_information]. The first group involves communication via some sort of

voluntary bodily motion named by the verb, including classes *crane-40.3.2*, *curtsey-40.3.3*, and *wink-40.3.1*. In addition to the basic **transfer_info** predicates, these classes take a Patient role that is shown to be a body part of the Agent with a **part_of**(Patient, Agent) predicate. During the course of the **transfer_info** subevent, the Agent moves the Patient into a verb-specific position, represented using **has_position**(e , Patient, V_Position) and **body_motion**(\ddot{e} , Agent). These classes have a more nuanced take on the possession boundaries than the basic representation in (8). In example (10) from *wink-40.3.1*, the Theme is a nonverbal emotional state conveyed through a bodily motion. We can generally assume that the Recipient does not have prior access to this type of information, and we make this explicit in e_1 .

(10) *Linda nodded her agreement*

Agent V Theme
has_information(E, Agent, Theme)
¬has_information(e_1 , ?Recipient, Theme)
¬has_position(e_1 , ?Patient, V_Position)
transfer_info(\ddot{e}_2 , Agent, Theme, ?Recipient)
body_motion(\ddot{e}_2 , Agent)
has_position(e_2 , ?Patient, V_Position)
has_information(e_3 , ?Recipient, Theme)
part_of(?Patient, Agent)
cause(\ddot{e}_2, e_3)

The second group involves potentially involuntary nonverbal expressions of an internal state, and includes classes *animal_sounds-38* and *nonverbal_expression-40.2* (11). As part of this release, we have added a new Stimulus thematic role to these classes. The previous release included frames for constructions using a Recipient, like *Paul laughed at Mary* and *The dog barked at the cat*, but didn't cover possible constructions like *Paul laughed at Mary to his friends* or *The dog whimpered to its owner about the rabbit in the yard*. Adding Stimulus and its usual predicate **in_reaction_to**(e , Stimulus) to these representations aligns them with the other Stimulus/Experiencer classes and expands the range of frames they cover. These classes reflect the same assumptions about boundaries on possession shown in (10).

(11) *The dog whimpered to its owner at the sight of the rabbit in the yard*

Agent	V	Recipient	Stimulus
has_information (E, Agent, ?Theme)			
-has_information (e_1 , Recipient, ?Theme)			
transfer_info (e_2 , Agent, ?Theme, Recipient)			
manner (e_2 , Agent, V_Manner)			
in_reaction_to (e_2 , Stimulus)			
has_information (e_3 , Recipient, ?Theme)			
cause (e_2 , e_3)			

5 Automatic VerbNet Parsing

To facilitate immediate use of the new VerbNet semantic representations, we are releasing a semantic parser that predicts the updated semantic representations from events in natural language input sentences. For a given predicative verb in a sentence, we define VerbNet semantic parsing as the task of identifying the VN class, associated thematic roles, and corresponding semantic representations linked to a frame within the class.

We approach VerbNet semantic parsing in three distinct steps: 1. Sense disambiguation to identify the appropriate VN class, 2. PropBank semantic role labeling (Gildea and Jurafsky, 2002; Palmer et al., 2005) to identify and classify arguments, and 3. Alignment of PropBank semantic roles with VN thematic roles within a frame belonging to the predicted VN class. After aligning arguments from the PropBank SRL system’s output with the thematic roles in a particular VN frame, the frame’s associated semantic predicates can be instantiated using the aligned arguments.

For sense disambiguation, we use a supervised verb sense classifier trained on updated VN class tags (Palmer et al., 2017). For semantic role labeling, we use a variation of the system described in He et al. (2017) and Peters et al. (2018) using solely ELMo embeddings (without any pre-trained or fine-tuned word-specific vectors) trained on a combination of three PropBank annotated corpora described in (O’Gorman et al., 2019): OntoNotes (Hovy et al., 2006), the English Web TreeBank (Bies et al., 2012), and the BOLT corpus (Garland et al., 2012). For alignment, we begin by applying updated SemLink mappings (Palmer, 2009) to map PropBank roles to linked VN thematic roles for the identified VN class. Remaining arguments are then mapped using heuristics based on the syntactic and selectional restrictions defined in the VN class. To se-

lect among multiple valid frames, we select the frame with highest total number of roles among the VN frames with the fewest unmapped roles.

This approach to VN parsing using multiple independent systems represents a simple baseline approach. We leave a more sophisticated, unified approach to VN semantic parsing to future work.

6 Conclusion

The fine-grained semantic representations presented here improve the consistency and precision of VerbNet’s verb semantics, offering a more useful modeling for the subevent structure of particular event types. This should improve VerbNet’s utility for human-robot and human-avatar interaction, and lend enhanced richness to applications aimed at temporal event sequencing.

All of the resources described in this paper are freely available. An online, browsable version of all the semantic representations is available through the Unified Verb Index at https://uvi.colorado.edu/uvi_search. A downloadable version can be accessed at https://uvi.colorado.edu/nlp_applications. A demo of the VerbNet Parser is at <http://verbnetsemantic-parser.appspot.com/>.

Acknowledgments

We gratefully acknowledge the support of DTRA 16-1-0002/Project 1553695, eTASC - Empirical Evidence for a Theoretical Approach to Semantic Components and DARPA 15-18-CwC-FP-032 Communicating with Computers, C3 Cognitively Coherent Human-Computer Communication (sub from UIUC) and Elementary Composable Ideas (ECI) Repository (sub from SIFT). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of DTRA or the U.S. government.

In addition, we thank our anonymous reviewers for their constructive feedback.

References

Omri Abend, Roi Reichart, and Ari Rappoport. 2008. A supervised algorithm for verb disambiguation into verbnet classes. In *Proceedings of the 22nd International Conference on Computational Linguistics-Volume 1*, pages 9–16. Association for Computational Linguistics.

- Norm Badler, Martha Palmer, and Rama Bindiganavle. 1999. Animation control for real-time virtual humans. *Communications of the ACM*, 42(7):65–73.
- Ann Bies, Justin Mott, Colin Warner, and Seth Kulick. 2012. English web treebank. *Linguistic Data Consortium, Philadelphia, PA*.
- Claire Bonial, Susan Windisch Brown, William Corvey, Martha Palmer, Volha Petukhova, and Harry Bunt. 2011. An exploratory comparison of thematic roles in verbnet and lirics. In *Workshop on Interoperable Semantic Annotation*, page 39.
- Johan Bos. 2008. Wide-coverage semantic analysis with boxer. In *Proceedings of the 2008 Conference on Semantics in Text Processing*, pages 277–286. Association for Computational Linguistics.
- Susan Windisch Brown, Dmitriy Dligach, and Martha Palmer. 2014. Verbnet class assignment as a wsd task. In *Computing Meaning*, pages 203–216. Springer.
- Susan Windisch Brown and Martha Palmer. 2012. Semantic annotation of metaphorical verbs with verbnet: A case study of climb and poison. In *Workshop on interoperable semantic annotation*, page 72.
- Susan Windisch Brown, James Pustejovsky, Annie Zaenen, and Martha Palmer. 2018. Integrating generative lexicon event structures into verbnet. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC-2018)*.
- Danilo Croce, Roberto Basili, Alessandro Moschitti, and Martha Palmer. 2012. Verb classification using distributional similarity in syntactic and semantic structures. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*, pages 263–272. Association for Computational Linguistics.
- Jennifer Garland, Stephanie Strassel, Safa Ismael, Zhiyi Song, and Haejoong Lee. 2012. Linguistic resources for genre-independent language technologies: user-generated content in bolt. In *Workshop Programme*, page 34.
- Daniel Gildea and Daniel Jurafsky. 2002. Automatic labeling of semantic roles. *Computational Linguistics*, 28(3):245–288.
- Ana-Maria Giuglea and Alessandro Moschitti. 2006. Semantic role labeling via framenet, verbnet and propbank. In *Proceedings of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the Association for Computational Linguistics*, pages 929–936. Association for Computational Linguistics.
- Adele E Goldberg. 2006. *Constructions at work: The nature of generalization in language*. Oxford University Press on Demand.
- Luheng He, Kenton Lee, Mike Lewis, and Luke Zettlemoyer. 2017. *Deep Semantic Role Labeling: What Works and What’s Next*. *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Eduard Hovy, Mitchell Marcus, Martha Palmer, Lance Ramshaw, and Ralph Weischedel. 2006. *OntoNotes: the 90% solution*. In *Proceedings of the human language technology conference of the NAACL, Companion Volume: Short Papers*, pages 57–60. Association for Computational Linguistics.
- Jena D Hwang. 2014. Identification and representation of caused motion constructions.
- Daisuke Kawahara and Martha Palmer. 2014. Single classifier approach for verb sense disambiguation based on generalized features. In *LREC*, pages 4210–4213.
- Karin Kipper, Anna Korhonen, Neville Ryant, and Martha Palmer. 2006. Extending verbnet with novel verb classes. In *Proceedings of LREC*, volume 2006, page 1. Citeseer.
- Karin Kipper and Martha Palmer. 2000. Representation of actions as an interlingua. In *Proceedings of NAACL-ANLP Workshop on Applied Interlinguas: practical applications of interlingual approaches to NLP*, pages 12–17. Association for Computational Linguistics.
- Beth Levin. 1993. *English Verb Classes and Alternations: A Preliminary Investigation*. University of Chicago Press.
- Edward Loper, Szu-Ting Yi, and Martha Palmer. 2007. Combining lexical resources: mapping between propbank and verbnet. In *Proceedings of the 7th International Workshop on Computational Linguistics, Tilburg, the Netherlands*.
- Anjali Narayan-Chen, Colin Graber, Mayukh Das, Md Rakibul Islam, Soham Dan, Sriraam Natarajan, Janardhan Rao Doppa, Julia Hockenmaier, Martha Palmer, and Dan Roth. 2017. Towards problem solving agents that communicate and learn. In *Proceedings of the First Workshop on Language Grounding for Robotics*, pages 95–103.
- Tim O’Gorman, Sameer Pradhan, James Gung, and Martha Palmer. 2019. The unified propbank landscape. *Ms*.
- Martha Palmer. 2009. Semlink: Linking propbank, verbnet and framenet. In *Proceedings of the generative lexicon conference*, pages 9–15. Pisa Italy.
- Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. The proposition bank: An annotated corpus of semantic roles. *Computational linguistics*, 31(1):71–106.

- Martha Palmer, James Gung, Claire Bonial, Jinho Choi, Orin Hargraves, Derek Palmer, and Kevin Stowe. 2017. The Pitfalls of Shortcuts: Tales from the word sense tagging trenches. *Springer series Text, Speech and Language Technology, Essays in Lexical Semantics and Computational Lexicography - In Honor of Adam Kilgarriff*.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237.
- J. Pustejovsky. 1995. *The Generative Lexicon*. Bradford Book. Mit Press.
- James Pustejovsky. 2013. Dynamic event structure and habitat theory. In *Proceedings of the 6th International Conference on Generative Approaches to the Lexicon (GL2013)*, pages 1–10. ACL.
- James Pustejovsky and Jessica Moszkowicz. 2011. The qualitative spatial dynamics of motion. *The Journal of Spatial Cognition and Computation*.
- James Pustejovsky, Martha Palmer, Annie Zaenen, and Susan Brown. 2016. Verb meaning in context: Integrating verbnet and gl predicative structures. In *Proceedings of the LREC 2016 Workshop: ISA-12, Potoroz, Slovenia*, volume 2016.
- J. Allbeck N. Badler A. Joshi R. Bindiganavale, W. Schuler and M. Palmer. 2000. Dynamically, altering agent behaviors using natural language instructions. In *Proceedings of the 4th International Conference on Autonomous Agents (AGENTS)*, pages 293–300.
- Karin Kipper Schuler. 2005. Verbnet: A broad-coverage, comprehensive verb lexicon.
- Lei Shi and Rada Mihalcea. 2005. Putting pieces together: Combining framenet, verbnet and wordnet for robust semantic parsing. In *International conference on intelligent text processing and computational linguistics*, pages 100–111. Springer.
- Annie Zaenen, Daniel G Bobrow, and Cleo Condoravdi. 2008. The encoding of lexical implications in verbnet predicates of change of locations. In *LREC*.