Mining and Organizing a Resource of State-changing Verbs

Derry Tanti Wijaya
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA, 15213
dwijaya@cs.cmu.edu

Ndapandula Nakashole
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA, 15213
ndapa@cs.cmu.edu

Tom M. Mitchell
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA, 15213
tom.mitchell@cs.cmu.edu

Abstract

Most knowledge bases (KBs) that emerged in recent years are static. They contain facts about the world yet are seldom updated as the world changes: attributes or possessions of entities in the KB change over time; what (actions) the entities can do or what can be done to them change over time. Current methods for extracting facts in the KBs do not yet detect such *state changes* although these changes have implications for the correctness of the KB. We propose a method for learning change-of-states in entities that are brought about by event-expressing verbs acting on the entities. These change of states in entities can be viewed as KB updates of facts pertaining to these entities. We envision a resource of *state changing verbs*, that contains *pre- and post-conditions* of each verb on its arguments: the entry condition of entities that allow an event expressed by the verb to take place and the condition of entities that will be true after the event occurs. When a state changing verb acts on an entity or when the entity's state change is detected in text, we can use the resource to update KB facts about the entity.

1 Introduction

Today's KBs are largely static [1, 2, 3]. They are seldom updated as the web changes when in reality new facts arise while others cease to be valid or change over time. One approach towards real-time population of KBs is to extract relational facts between entities from dynamic content of the web such as news, blogs and user comments in social media [4]. However, current methods for extracting relational facts do not differentiate between patterns for extracting facts with those pertaining to state changes. For example, although the verbs: *marry* and *divorce* are both good patterns for extracting HASSPOUSE relational facts, they cause different change of states in their arguments – one causes its arguments to change state from "single" to "married" while the other from "married" to "single". In other words, their arguments have different sets of *pre-* and *post-conditions*. These pre- and post-conditions can be useful for differentiating that "marry" signals the beginning of HASSPOUSE relation while "divorce" signals the end – a useful information for updating this KB relation.

This paper proposes recasting KB updates as a new task of identifying state changes of entities brought upon by state-changing verbs. To realize such an approach, we propose to construct a resource of *state changing verbs*. This includes the *pre- and post-conditions* of each verb on its arguments: the entry condition of entities that allow an event expressed by the verb to take place and the condition of entities that will be true after the event occurs.

Building such a resource and recasting KB updates as state change detection, bring about several other advantages:

1. Even when relational facts between entities are not stated explicitly in text (e.g., via relational verb phrases), relations can still be inferred from change of state in entities. For

example, if an entity changes state from being *single* to being *married*, then we can infer that there is a HASSPOUSE relation existing between the entity and another entity in text.

2. Typical current state of entities can be inferred from state changing verbs acting on them. For example, if there is a state changing verb *fire* acting on an entity, then we can infer that the entity is typically now *unemployed*.
3. Detecting when a state change occurs to an entity can used for temporally scoping the corresponding KB facts about the entity [5].
4. Useful event sequences such as scripts [6] can be modeled as a collection of events chained together by the pre-condition and post-condition overlap of their shared entities.
5. Knowledge of pre-condition and post-condition of entities opens up interesting opportunities for other inferences and reasoning e.g., to infer links between semantic roles of related events, or to infer how the effect of one event can be *cascaded* down via the pre- and post-condition of the shared entities to other entities.

Verb Resource

```
marry(Partner_1, Partner_2, time_marry)
  pre-condition:
    adjective(Partner_1) = [single, engaged]
    adjective(Partner_2) = [single, engaged]
  post-condition:
    adjective(Partner_1) = [married]
    adjective(Partner_2) = [married]
  is-typically-temporally-before:
    divorce(Partner_1, Partner_2)
    bear(Mother, Father, Child)
  is-typically-temporally-after:
    propose(Speaker, Addressee)
      (Speaker = Partner_1 && Addressee = Partner_2)
      | (Speaker = Partner_2 && Addressee = Partner_1)
  expresses-relation:
    HASSPOUSE(Partner_1, Partner_2,
              begin_time, end_time)
      begin_time = time_marry
```

```
bear(Mother, Child, time_bear)
  pre-condition:
    adjective(Mother) = [pregnant, expecting]

  post-condition:
    verb(Mother) = [give_birth]
    adjective(Child) = [born]
  is-typically-temporally-before:
    study(Student, Institution)
     (Student = Child)
  is-typically-temporally-after:
    marry(Partner_1, Partner_2)
     (Partner_1 = Mother if ISAFEMALE(Partner_1))
      | (Partner_2 = Mother if ISAFEMALE(Partner_2))
  expresses-relation:
    HASCHILD(Mother, Child,
              begin_time, end_time)
      begin_time = time_bear
```

Figure 1: An illustration of the proposed resource of state-changing verbs.

2 Proposed Resource

Figure 1 is an illustration of the proposed verb resource. For each verb, when available, the pre- and post-condition of its arguments (semantic roles) are listed. The pre- and post-condition of each argument are word vectors containing the frequency of verbs, adjectives, and nouns co-occurring with an argument that express the state of the argument before and after it participates in the event expressed by the verb. These verbs, adjectives and nouns in the vectors intuitively represent what the argument does/what is done to it i.e., the verbs for which the argument is the *Subject/Object*, how it is described (i.e., its attributes), and what it possesses.

For example, the pre-condition of *Partner_1* of the verb *marry* is a vector containing the adjectives *single* and *engaged*, while the post-condition is a vector containing the adjective *married*. Note that some post-condition vectors can contain the words in the pre-condition vectors but with a decrease in frequency. For example, the post-condition of *marry* can contain the word *single* at a decreased frequency. The change in frequency can be used to detect incidence of state change.

For each verb in the resource, whenever available, verbs that typically follow/precede it are also listed. These can be learned from the corpora (using narrative script finding [7] or textual entailment methods [8]) or inferred from the pre- and post-condition of verbs. For example, the verbs *divorce* and *propose* can be inferred to follow/precede the verb *marry* from the pre- and post-condition of their arguments while other verbs like: *bear* are verbs that often follow/precede *marry* in the corpora without requiring that their arguments' pre-/post-condition match *marry*'s post-/pre-condition.This temporal sequence is another useful resource about verbs. Temporal sequences of verbs can be useful for predicting temporal sequences of events, for example. They can also be useful to better disambiguate the meaning of a verb in a sentence, given the presence of other verbs in the document.

Whenever verbs are temporally linked, the links between their semantic roles are also added. These links can be useful for inferring semantic role labels or for resolving co-reference involving temporally linked verbs, for example. The mapping to the KB relation is also included. For example, the verb *marry* is mapped to HASSPOUSE relation and the time of *marry* is mapped to the **begin_time** of HASSPOUSE i.e., the verb *marry* causes a state change signaling the beginning (i.e., the creation) of a HASSPOUSE fact over entities expressed by *marry*'s arguments. This mapping to the KB facilitates knowledge base updates of an entity's facts when a state changing verb acts on the entity.

3 Background

3.1 Theory of Verbs

State changing verbs have their background in linguistics. In the linguistic study of verb meaning, verbs can be categorized into stative or dynamic [9]. The stative verbs express states of the arguments while the dynamic verbs express events (e.g., *want* vs. *acquire*, *own* vs. *buy*, *work* vs. *hire*). In particular, *state changing* verbs, also known as change-of-state verbs in linguistics, express the change of state of the arguments participating in the event [10]. Stative verbs have a simple event structure: [y ⟨**STATE**⟩] or [**BECOME** [y ⟨**STATE**⟩]] while state changing verbs have a complex event structure: [[x **ACT**] **CAUSE** [**BECOME** [y ⟨**STATE**⟩]]] where x and y are verbs' arguments.

Although state changing verbs have been studied extensively in linguistics, there is not yet a lexical resource that compiles the actual change of state brought about by these verbs on their arguments. A lexical resource that represents the meanings of these verbs as changes in the states of their arguments is compelling because it opens a range of new inference options for KB updates and more generally for natural language understanding.

3.2 Conceptual Dependency Theory

The idea of representing events as concepts (*objects* and *actions*) and the change-of-state caused by the *actions* on the *objects* can be found in Conceptual Dependency Theory [6]. This theory focuses on extracting concepts from the sentence and modeling their dependencies: *objects* can perform *actions*, *actions* can be performed on *objects*, *objects* can be described by *states*, and *actions* can change the state of *objects*.

The advantage of Conceptual Dependency Theory is that it is a language-independent meaning representation. Words can trigger conceptual dependency structures that can provide predictions about what will come next. It also facilitates inference – unstated change-of-state or states of unknown words can be inferred. Inference is also attached to concepts so there is no need for complex inference rules. However, the critique of this theory is that it is incomplete. There is no ontology of *objects* or *actions*. The *states* of the objects are defined ad-hoc (e.g., the state of being "dead" means having a "health" score of -10 while being "tip top" means having a score of +7). The set of chosen *actions* are also critiqued in terms of whether it is sufficient to represent all the possible events in the world. To resolve these issues, we propose to focus instead on verbs for expressing *actions*. This will provide a high-coverage *vocabulary* of events described in natural language. Secondly, by relating *objects* and *actions* to *entities* and *relations* in knowledge bases, we can give an ontology to *objects* and *actions* and a clear definition of *states* and change-of-states as relation instances and their updates in the knowledge base.

3.3 Pre- and Post-condition from Time-Stamped Corpora

In our previous work [11], we discovered that we can automatically identify changes that occur to an entity based on the changes in the words surrounding the entity over time. By clustering the words that surround the entity over time, we identify *when* (at which year) changes occur, and also *what* changes occur (i.e., what clusters of words are in transition).

We learn two important insights from this work: (1) that some changes occurring to an entity can be identified from the changes in the words surrounding the entity over time and (2) that some changes occurring to an entity can coincide with events happening to the entity. We want to extend this work by learning for *state changing* verbs (i.e., verbs that express events *and* change the state of their arguments), the change brought about by these verbs on the arguments. We propose to learn this change from the change in the words surrounding entities expressed by these arguments over time.

4 Sources of Pre-and-post Conditions of Verbs

The first challenge to building a resource on the pre- and post-conditions of verbs is sparsity in the source documents from which to extract these pre- and post-conditions. One source may not be enough for learning pre- and post-condition of verbs. Some pre- and post-condition may not always be mentioned explicitly in sentences, or they may only be mentioned long before or long after the event that causes the state change happened.

Therefore, we propose to restrict ourselves to learning *instantaneous* state change – state change that happens instantly after the event happens. We also propose to utilize signals from several corpora of different nature such that the sparsity of one may be compensated by the redundancy of another.

4.1 Corpus Statistics

Inspired by the previous work [11], we can leverage the idea that the change happening to entities participating in a state changing event can be identified from the change in the words surrounding the entities over time, for example results, see Table 1. In this table, we can see that the word uni-grams and bi-grams before and after the events expressing the relation PRESIDENTOF(*Person*, *US*) and WINAWARD(*Movie*, *Best Picture*) reflect meaningful events and state changes happening to the entities. These context show the potential for learning typical pre-and-post conditions of state-changing verbs for the PRESIDENTOF relation such as *elect* and the state-changing verbs for the WINAWARD relation such as *win*.

Since one corpus may be sparse for learning pre- and post-condition of verbs, we can make use of corpora of different time-stamp granularity to complement one another. For example, the effect of the verb "marry" (e.g., "married", "spouse") may not be mentioned in the news document the day after the marriage event happened but they may be mentioned in the biographies of the entities in books. Hence, the sparsity of news documents with the day granularity (GigaWord [12]) can be overcome by the redundancy of documents with the year granularity (Google Books N-gram [13]).

4.2 Wikipedia Edit History

Another challenge to learning the pre- and post-conditions of verbs is that they may not always be mentioned explicitly in text. Some pre-condition may not be mentioned explicitly but *edited* out of the document when the post-condition occurs. Consider the state changing verb, "married". its typical pre-condition "engaged" is often omitted from documents that talk about the entity that is now no longer just engaged but married. This pre-condition can be extracted, for example, from Wikipedia edit history of the relevant entities. Methods to automatically classify Wikipedia edits into their different types [14] can be used to filter Wikipedia edits for only those relevant to the state change in entities.

4.3 Lexical Resources

Another challenge to learning the pre- and post-conditions of verbs is that for some state-changing verbs their pre- and post-conditions are such basic knowledge that they are rarely mentioned in text e.g., *sadden*, *anger*, etc. Some state-changing verbs are also very rare that the verbs themselves are rarely mentioned in general purpose text e.g., *achromatize*.

Therefore, we propose to restrict ourselves and start only with state-changing verbs that can be found aplenty in general purpose corpora such as GigaWord [12], Google Books N-gram [13], Wikipedia, and ClueWeb [15]. For the purpose of knowledge base updates, we will also focus first on verbs that can express existing relations in the knowledge bases such as YAGO [1], NELL [2] and KBP [1].

[1]http://www.nist.gov/tac/2012/KBP/index.html

Relation	Context Before	Context After
US President	was elected, took office, became president, vice president, senator, governor, candidate	by president, administration,
Best Picture	hour minute, star, nominated for, to win, won the, was nominated	best picture, academy award, oscar, nominated, won, best actress, best actor, best supporting

Table 1: Example of various contexts learned automatically from a time-stamped corpus, before and after the event expressing the relation happens. These context show the potential for learning typical pre-and-post conditions of events.

In future, for verbs whose pre- and post-condition are basic knowledge, we can start from dictionary definitions of verbs to extract their pre- and post-condition. For example, from the WordNet [16] definition of the verb *anger* as "become angry", we can infer that the post-condition of this verb *anger* contains adjectives such as "angry".

Lexical resources of verbs such as VerbNet [17] also contains useful diagnostics for detecting pre- and post-condition of verbs from the parse information and the syntactic realizations of verbs and their semantic roles in sentences. For example, the verb *deport* can appear syntactically in sentences as "Agent *deport* Theme *to* Destination" which translates semantically in VerbNet to CAUSE(Agent, Event), LOCATION(START(Event), Theme, ?Source), LOCATION(END(Event), Theme, Destination). We can therefore infer that *deport* causes LOCATION changes in its pre- and post-condition.

Other lexical resource such as FrameNet that contains sequences of causally related frames can be used to infer pre- and post-condition of verbs that evoke these frames. The *pre*-condition of some verbs that evoke a frame can be constrained by or inferred from the *post*-condition of verbs that evoke a causally-preceding frame. For example, the *pre*-condition of the verb *investigate* that evokes the frame CRIMINALINVESTIGATION can be constrained by or inferred from the *post*-condition of the verb *perpetrate* that evokes the frame COMMITTINGCRIME, a causally-preceding frame. An interesting research direction will be to integrate signals from all these different sources to come up with the overall pre- and post-condition of verbs.

Another challenge to learning pre- and post-condition of verbs is that the pre- and post-conditions of a state changing verb may depend also on the argument of the verb or its type. For example, the verb *elect* causes different change of states when applied to different argument types: a "president" vs. a "vice president", for example. Another example, the verb *schedule* will have different pre- and post-condition when applied to different argument types: a "flight" vs. a "meeting". To solve this, we propose to organize our verb resource according to the argument types of the verbs. We will therefore learn the pre- and post-conditions of typed-verbs – verbs whose arguments (semantic roles) are typed. Hence, the verb *schedule* will have a couple of entries in the resource, each representing the different argument types it can take. The different entries of the verb in the resource can be seen as representing the different "senses" of the verb.

Furthermore, since lexical pre- and post-condition of verbs can vary widely, we can generalize the lexically varied pre- and post-condition by mapping them to categories and therefore relation values in the KB. For example, the pre- and post-condition of the verb *elect* that contain nouns such as "president" and "vice president"; these nouns can be mapped to the values of the HASJOBPOSITION relation in the KB. We can then generalize that the verb *elect* causes a change in the HASJOBPOSITION relation.

5 Conclusions

We envision a state changing verbs resource that contains pre- and post-condition of each verb on its entities. We have already demonstrated how useful such a resource would be for temporal scoping of relation instances (facts) in the KB [5]. Such a resource can enrich other tasks related to knowledge updates and natural language understanding. In this work we focus on building a resource on state-changing verbs. However, other types of words such as nouns can cause state change in entities e.g., “Microsoft’s *acquisition* of Skype”. In future, it will be interesting to explore if we can extend the resource on state-changing verbs to more generally, state-changing frames, for example.

References

[1] Fabian M Suchanek, Gjergji Kasneci, and Gerhard Weikum. Yago: a core of semantic knowledge. In *Proceedings of the 16th international conference on World Wide Web*, pages 697–706. ACM, 2007.

[2] Andrew Carlson, Justin Betteridge, Bryan Kisiel, Burr Settles, Estevam R Hruschka Jr, and Tom M Mitchell. Toward an architecture for never-ending language learning. In *AAAI*, volume 5, page 3, 2010.

[3] Anthony Fader, Stephen Soderland, and Oren Etzioni. Identifying relations for open information extraction. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 1535–1545. Association for Computational Linguistics, 2011.

[4] Ndapandula Nakashole and Gerhard Weikum. Real-time population of knowledge bases: opportunities and challenges. In *Proceedings of the Joint Workshop on Automatic Knowledge Base Construction and Web-scale Knowledge Extraction*, pages 41–45. Association for Computational Linguistics, 2012.

[5] Derry Wijaya, Ndapa Nakashole, and Tom Mitchell. Ctps: Contextual temporal profiles for time scoping facts via entity state change detection. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2014.

[6] Roger C Schank and Robert P Abelson. *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures*. Psychology Press, 2013.

[7] Nathanael Chambers and Daniel Jurafsky. Unsupervised learning of narrative event chains. In *ACL*, pages 789–797. Citeseer, 2008.

[8] Hila Weisman, Jonathan Berant, Idan Szpektor, and Ido Dagan. Learning verb inference rules from linguistically-motivated evidence. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 194–204. Association for Computational Linguistics, 2012.

[9] David R Dowty. *Word meaning and Montague grammar: The semantics of verbs and times in generative semantics and in Montague’s PTQ*, volume 7. Springer, 1979.

[10] Charles John Fillmore. *The grammar of’hitting’and’breaking’*. 1968.

[11] Derry Tanti Wijaya and Reyyan Yeniterzi. Understanding semantic change of words over centuries. In *Proceedings of the 2011 international workshop on DETecting and Exploiting Cultural diversiTy on the social web*, pages 35–40. ACM, 2011.

[12] David Graff, Junbo Kong, Ke Chen, and Kazuaki Maeda. English gigaword. *Linguistic Data Consortium, Philadelphia*, 2003.

[13] Jean-Baptiste Michel, Yuan Kui Shen, Aviva Presser Aiden, Adrian Veres, Matthew K Gray, Joseph P Pickett, Dale Hoiberg, Dan Clancy, Peter Norvig, Jon Orwant, et al. Quantitative analysis of culture using millions of digitized books. *science*, 331(6014):176–182, 2011.

[14] Johannes Daxenberger and Iryna Gurevych. Automatically classifying edit categories in wikipedia revisions. In *EMNLP*, pages 578–589, 2013.

[15] J. Callan, M. Hoy, C. Yoo, and L. Zhao. Clueweb09 data set. *boston.lti.cs.cmu.edu*, 2009.

[16] George A Miller, Richard Beckwith, Christiane Fellbaum, Derek Gross, and Katherine J Miller. Introduction to wordnet: An on-line lexical database*. *International journal of lexicography*, 3(4):235–244, 1990.

[17] Karin Kipper, Hoa Trang Dang, and Martha Palmer. Class-based construction of a verb lexicon. In *AAAI/IAAI*, pages 691–696, 2000.