

The Predicate Matrix and the Event and Implied Situation Ontology: Making More of Events

Roxane Segers,¹ Egoitz Laparra,² Marco Rospocher,³ Piek Vossen,¹ German Rigau,², Filip Ilievski¹

¹The Network Institute, Vrije Universiteit Amsterdam

²IXA Group, UPV/EHU

³Fondazione Bruno Kessler

r.h.segers@vu.nl, egoitz.laparra@ehu.es, rosposcher@fbk.eu, piek.vossen@vu.nl,
german.rigau@ehu.es, f.ilievski@vu.nl

Abstract

This paper presents the Event and Implied Situation Ontology (ESO), a resource which formalizes the pre and post situations of events and the roles of the entities affected by an event. The ontology reuses and maps across existing resources such as WordNet, SUMO, VerbNet, PropBank and FrameNet. We describe how ESO is injected into a new version of the Predicate Matrix and illustrate how these resources are used to detect information in large document collections that otherwise would have remained implicit. The model targets interpretations of situations rather than the semantics of verbs per se. The event is interpreted as a situation using RDF taking all event components into account. Hence, the ontology and the linked resources need to be considered from the perspective of this interpretation model.

1 Introduction

In this paper, we present the new release of the Event and Implied Situation Ontology (ESO) that is matched with a new version of the Predicate Matrix (PM). Both resources rely on Semantic Role Labeling (SRL) descriptions and are used to detect and abstract over events, their participants and event implications in a large document collection about ten years of global automotive industries.

ESO (Segers et al., 2015) is a newly developed domain ontology to enhance the extraction and linking of dynamic and static events and their implications in text. Explicit modeling of event implications allows for extracting sequences of states and changes over time regardless of if this information was directly expressed in text, or inferred

by a reasoner. Figure 1 shows such a chain of expressions for dynamic (*hire, starts at, fire, leave*) and static events (*works for, employs, is CEO*) and their implied situations. Lexicons that define implications of events, e.g. VerbNet (Kipper et al., 2000; Kipper et al., 2006), are rare and usually focus on the meaning of verbs in isolation. However, lexical structures do not make explicit how the meaning of a verb needs to be combined with other event components, such as the participants and the temporal properties for the purpose of semantic parsing. We therefore follow an ontological approach to interpret situations on the basis of text interpretation of all the event components to make the implications explicit. Though some research on deductive reasoning over Frame annotated text (e.g. (Scheffczyk et al., 2006)) and defining pre and post situations of predicates exist (Im and Pustejovsky, 2009; Im and Pustejovsky, 2010), to the best of our knowledge, ontologies that model both events, roles and implications do not. Most closest comes the extension to DOLCE-LITE (Hicks, 2009) that models property values as quality regions for reasoning. However, these quality regions are not connected to the events in the ontology as pre and post situations. Axioms in generic and top ontologies such as SUMO (Niles and Pease, 2001) and DOLCE (Masolo et al., 2002) provide a comprehensive semantic specification of the concepts, but these axioms do not always provide the information relevant and specific for our domain. Furthermore, such ontologies need to be integrated with semantic parsing systems that deal with expressions on natural language to be able to test these models. We therefore decided to develop a new ontology for modeling static and dynamic events and their implications that is tailored to a semantic parsing system for text.

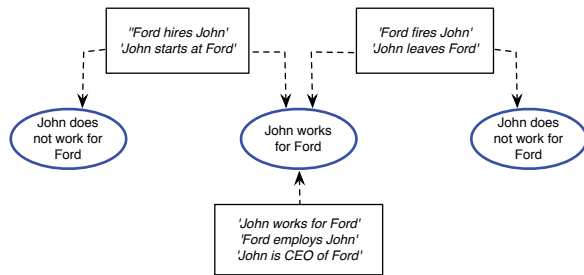


Figure 1: A chain of dynamic and static event expressions and their implied situation.

Version 2 of ESO was released in July 2015 and now includes modeling of scalar values and an extended expressivity of the assertions that define the situation that holds before, during and after the event. It also includes updated mappings to SUMO classes and to FrameNet frames and Frame Entities.

The Predicate Matrix¹ (de Lacalle et al., 2014a; de Lacalle et al., 2014b) is the second resource presented in this paper. It integrates predicate and role information from e.g. FrameNet, VerbNet, PropBank, NomBank and WordNet. This resource is used to assign role and predicate annotations at sentence level. All classes and roles in ESO are fed back into to the Predicate Matrix. As such the ontology provides an additional layer of annotations in text that allow for inferencing over events and implications. The current version of the PredicateMatrix contains 8,495 predicates from PropBank and NomBank connected to 4,704 synsets of WordNet, 554 frames of FrameNet and 55 ESO classes. On the other hand, this resource contains 23,386 roles of PropBank and NomBank mapped to 2,343 frame-elements of FrameNet and 53 ESO roles.

The remainder of this paper is organized as follows. Section 2 presents the ontological meta model and the content of ESO. Section 3 describes the Predicate Matrix and the integration with ESO. In section 4 we provide a preliminary overview of the Predicate Matrix and ESO in our document collection. In section 5 we report on an experiment carried out on a small corpus. We conclude in section 6 with a discussion and some outlines for future work.

¹<http://adimen.si.ehu.es/web/PredicateMatrix>

2 The Event and Implied Situation Ontology

In this paragraph, we briefly describe the meta model of the ontology with a focus on the instantiation of the rules that define what situation holds before, during and after some event. Next, we describe how the ontology was built and we provide an overview of the current content. The ESO ontology and a detailed documentation can be found online: <https://github.com/newsreader/eso>

ESO is an OWL 2 ontology.² It assumes that the semantic representation of text is converted to an RDF representation of event and entity instances, between which relations are expressed as triples. For instance, the statement

```
:obj-graph-eventX {
  :eventX
    a          eso:Translocation;
    eso:translocation-theme  :Enzo Ferrari;
    eso:translocation-goal   :Rome;
    eso:translocation-source :Napels;
    sem:hasTime  :time_eventX.
}
```

specifies that the event (X) is of a certain type (eso:Translocation), that it involves an entity playing the role of a moving thing (:Enzo Ferrari), an entity playing the role of goal (:Rome), an entity playing the role of source (:Napels) and that it occurred at a certain time (:time_eventX). From these representations, we derive the statements that express the pre, post and during event situations.

For this purpose, we defined five core classes in ESO: 1) **Event**: this class is the root of the taxonomy of event types. Any event detected in a text will be an instance of some class of this taxonomy; 2) **DynamicEvent**: this is a subclass of Event for which dynamic changes are defined; 3) **StaticEvent**: this is another subclass of Event for “static” event types which capture more stable circumstances; 4) **Situation**: the individuals of this class are actual pre, post and during situations that will be instantiated starting from the event instances detected in the text; 5) **SituationRule**: the individuals of this class enable to encode the rules for instantiating pre/post/during situations when a certain type of event is detected.

Further, ESO includes mapping properties to match ESO roles to FrameNet roles, and properties to match ESO classes to FrameNet frames and SUMO classes. The mappings to FrameNet are necessary to translate the annotations provided by the SRL module using the Predicate Matrix to our

²<http://www.w3.org/2001/sw/wiki/OWL>

ontology. This is then exploited by the reasoning module to instantiate situations from events.

2.1 Formalization of the rules for instantiating situations from events

For all event classes in ESO an `eso:SituationRule` is defined; the individuals of this class trigger the pre, post and during situation related to a class or a set of event classes. For instance, the class `eso:Translocation` has two specific individuals: `pre_Translocation` and `post_Translocation`. Each `eso:SituationRule` individual defines exactly how the triples inside the Situation named graph have to be defined. This is done by defining an individual for each assertion to be created, which has three annotation properties: `eso:hasSituationAssertionSubject` (a role to be used as subject in the assertion), `eso:hasSituationAssertionObject` (a role to be used as object in the assertion) and `eso:hasSituationAssertionProperty` (a property relating the subject and object). In the case of e.g. `eso:Translocation`, the individual `pre_Translocation` has two `eso:SituationRuleAssertions`, where e.g. `eso:pre_Translocation_assertion_1` states:

```
eso:pre_Translocation_assertion1
  eso:hasSituationAssertionSubject    eso:translocation-theme;
  eso:hasSituationAssertionProperty    eso:atPlace;
  eso:hasSituationAssertionObject      eso:translocation-source.
```

Based on all class assertions, the ESO reasoner³ can now infer that some event belongs to the class `eso:Translocation` and that it has entity instances in certain roles where some entity is at some place before the event and not at this place after the event. The instantiation of the defined situations for the example event instance of `eso:Translocation` will then look as follows:

```
:eventX_pre {
  :Enzo Ferrari      eso:atPlace      :Napels
  :Enzo Ferrari      eso:notAtPlace   :Rome
}:eventX_post
  :Enzo Ferrari      eso:atPlace      :Rome
  :Enzo ferrari      eso:notAtPlace   :Napels
}
```

Instantiation of events that express a change in a scalar value By default, situation assertions will only fire if some instance for an ESO role is found by the SRL module. However, in specific cases we also allow that assertions are instantiated even though no instance exists for the ESO role. We do this by adding an OWL existential restriction on the event class for the role considered. The reasoner will check if an instance of

³Implemented as a processor of RDFpro (Corcoglioniti et al., 2015b). See also: <http://bit.ly/ESOreasoner>

the role exists, if not it will create a blank node. This OWL existential restriction is applied in ESO for event classes that express a relative change in the value of an attribute (e.g. `eso:Damaging`, `eso:Increasing`, `eso:Attacking`) where the attribute itself such as 'price' or 'damagedness' often remains implicit. As such, it is possible to assert statements based on 'incomplete' information if needed. For `eso:Increasing`, the existential restriction is defined as follows:

```
eso:Increasing rdfs:subClassOf [
  a owl:Restriction ;
  owl:onProperty    eso:triggersPreSituationRule ;
  owl:hasValue      eso:pre_Increasing ] .
eso:Increasing rdfs:subClassOf [
  a owl:Restriction ;
  owl:onProperty    eso:triggersPostSituationRule ;
  owl:hasValue      eso:post_Increasing ] .
eso:Increasing rdfs:subClassOf [
  a owl:Restriction ;
  owl:onProperty    eso:quantity-attribute ;
  owl:someValuesFrom owl:Thing ] .

eso:pre_Increasing a eso:SituationRule .
eso:post_Increasing a eso:SituationRule .
```

These are the situation rule assertions defined for the pre and post situation of `eso:Increasing`:

```
eso:pre_Increasing_assertion1
  eso:hasSituationAssertionSubject    eso:quantity-item;
  eso:hasSituationAssertionProperty    eso:hasAttribute;
  eso:hasSituationAssertionObject      eso:quantity-attribute.

eso:pre_Increasing_assertion2
  eso:hasSituationAssertionSubject    eso:quantity-attribute;
  eso:hasSituationAssertionProperty    eso:hasRelativeValue;
  eso:hasSituationAssertionObjectValue '-'

eso:post_Increasing_assertion1
  eso:hasSituationAssertionSubject    eso:quantity-item;
  eso:hasSituationAssertionProperty    eso:hasAttribute;
  eso:hasSituationAssertionObject      eso:quantity-attribute.

eso:post_Increasing_assertion2
  eso:hasSituationAssertionSubject    eso:quantity-attribute;
  eso:hasSituationAssertionProperty    eso:hasRelativeValue;
  eso:hasSituationAssertionObjectValue '+'
```

The pre and post situation named graphs for the example sentence "Ford increased the production" can now be instantiated as follows:

```
:eventX_pre {
  :production      eso:hasAttribute      :xyz123
  :xyz123          eso:hasRelativeValue   '- '
}:eventX_post
  :production      eso:hasAttribute      :xyz123
  :xyz123          eso:hasRelativeValue   '+ '
}
```

These instantiations can be paraphrased as follows: the production has some unknown attribute and the value of this attribute has become more (+) after the event then it was before the event (-), meaning that the production goes from less (-) to more (+).

Alternatively, if the attribute is known, the assertions will instantiate the role that models the actual attribute. For a sentence like "Ford increased the price of the components", the event will look as follows:

```
:eventX_pre a eso:Increasing ;
eso:quantity-item      :component ;
eso:quantity-attribute :price ;
```

and the assertions will be instantiated as:

```

:eventX_pre {
  :component      eso:hasAttribute      :price
                  eso:hasRelativeValue  '- '
  :price
:eventX_post
  :component      eso:hasAttribute      :price
                  eso:hasRelativeValue  '+ '
  :price
}

```

Even though it may appear that these assertions for relative values are superfluous, we argue that finding *multiple* mentions of such an event and assertions over time, either with or without explicit values and attributes, allows for estimating the fluctuation of a certain value and the speed of the value change. We also need these values to determine that different event descriptions are coreferential even if one does not make the value explicit, while the other does. An existential representation of a value thus can match with an explicit value but two different explicit values cannot.

2.2 Mappings from external resources to ESO

A key ingredient of the ESO ontology is the mapping of FrameNet frames and Frame elements to the event types and roles that we defined. This mapping is necessary to translate the role annotations provided by the SRL module to our ontology vocabulary, which is then exploited by the reasoning module to instantiate situations from events. For each ESO event class and each ESO role we defined mapping properties representing the corresponding frames and frame elements. For instance, `eso:Giving` has three mappings to the frames `fn:Giving`, `fn:Sending`, and `fn:Supply`, meaning that if a frame of type `fn:Supply` or any of the others is identified in the text, it has to be considered as an event of type `eso:Giving`, and therefore pre and post situation rules defined for `eso:Giving` should be triggered. Similarly, the role `eso:possession-owner.1` is mapped to a set of frame elements. These mappings make clear that our ontology is providing only a partial definition for concepts. We only define those elements necessary for capturing salient pre and post situations of events and not any other meaning aspect. As such the implications of a change in ownership of something are similar for all instances of `eso:ChangeOfPossession`, such as *stealing*, *giving* or *seizing*.

2.3 Development and content of ESO (Vers. 2)

Version 2 of ESO was released in July 2015. It contains: a hierarchy of event classes; a set of

properties for the defining the pre, post and during situations of an event, and a set of roles for the entities affected by an event. In this section, we report how these structures were built and we conclude with an overview of its content.

The ESO ontology is a hand-built resource, based on high-frequent FrameNet frames that were extracted from a large domain-specific document collection. Frames that denote events pertaining to communication, feelings and perception were not taken into account. For deriving an initial conceptual structure for the frequent frames, we decided to map the frames manually to the SUMO ontology⁴ as a background model was based on the fact that it is freely available, well-documented, has a good coverage and is mapped to English WordNet and also the Predicate Matrix. As such, we derived four main conceptual clusters that formed the backbone of ESO: 'changes in possession', 'translocations', 'internal changes' and 'intentional events'. Next, we modeled 103 FrameNet frames into 63 distinct ESO event classes. Frames that denote fine-grained semantic distinctions are often grouped into one class in ESO since these distinctions do not influence the modeling of a salient set of pre and post situations. As such we build an event class hierarchy that reuses and maps to groups of FrameNet frames, which deviates from the approach taken in e.g. (A. Nuzozese, 2012) where FrameNet frames and frame relations are converted to RDF and partly to OWL. The second and third component of the ontology consists of properties and roles which are used for defining the assertions of the pre, post and during situations. All properties are hand-built, based on the shared semantics of the predicates related to a FrameNet frame and ESO class. The ESO roles define what entities are affected by a change and serve as the domain and range of properties. The majority of the ESO roles is mapped to a selection of FrameNet Frame Elements (FEs); these were selected manually from the FrameNet frames that correspond to an ESO class. This implies that not all FEs of a frame are mapped to ESO but only those that play a role in the assertions.

An important modeling decision is that assertions are defined at the highest possible level in the ontology. This way, all subclasses will inherit the same assertions and roles, which reduces re-

⁴<http://www.ontologyportal.org>

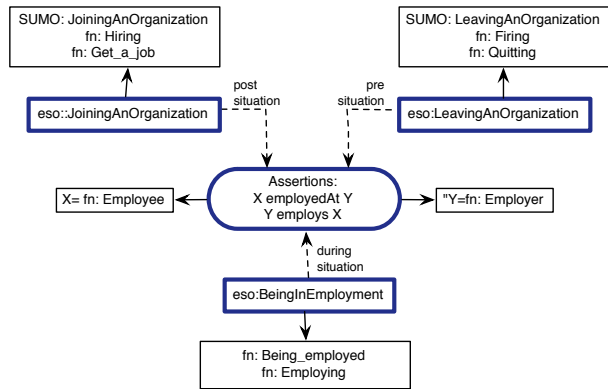


Figure 2: The shared assertion properties of a static and a dynamic event

dundancy. As such, many ESO roles have mappings to FEs that are aggregated from all mappings from ESO classes to FrameNet frame in a given sub-hierarchy. Another notable modeling choice is that the assertion properties for static event classes are partially shared with the assertion properties of the dynamic event classes. This is illustrated in Figure 2. Here, the same properties (`eso:employedAt` and `eso:employs`) are used in the pre situation assertion for the dynamic event class `eso:LeavingAnOrganization`, in the post situation assertion for the dynamic event class `eso:JoiningAnOrganization` and in the during situation assertion of the static event class `eso:BeingInEmployment`. As a result, the relation between the inferred situation of a dynamic event and the explicit mention of some state by a static event becomes explicit. Modeling the properties this way facilitates querying for chains of related changes and states (See also section 5).

To illustrate the expressivity of the assertions, in figure 3 we provide a non-formal transcription of a typical class in ESO, including the class mappings to SUMO and FrameNet, the aggregated role mappings to FEs, the inherited and class specific situation assertions and an example of the instantiation. From the "knowledge" in the example sentence, we are able to infer that a) Marie has 600 dollar and not the car before the event, while John does have the car but not the 600 dollar, b) after the event, the money and the car have changed ownership while c) the car itself has a value of 600 dollar during the exchange.

In table 1 we provide an overview of the content of ESO, including the number of mappings to FrameNet frames (103), SUMO classes (46) and

-FinancialTransaction: subclassOf: ChangeOfPossession
 "The subclass of ChangeOfPossession where some item changes of ownership in exchange for money."

Class mappings:
 closeMatch: fn:CommercialTransaction
 closeMatch: sumo:FinancialTransaction

Role mappings:
 possession-financial-asset: fn:Money

Inherited role mappings:
 possession-owner_1: fn:Supplier, fn:Exporter, fn:Donor, fn:Victim, fn:Source, fn:Lender, fn:Exporting_area, fn:Sender, fn:Seller
 possession-owner_2: fn:Perpetrator, fn:Importing_area, fn:Importer, fn:Lessee, fn:Buyer, fn:Recipient, fn:Borrower, fn:Agent
 possession-theme: fn:Theme, fn:Goods, fn:Possession
 possession-financial-asset: fn:Money

Assertions:

pre situation	possession-owner_1	notHasInPossession	poss.-financial-asset
	possession-owner_2	hasInPossession	poss.-financial-asset
post situation	possession-owner_1	hasInPossession	poss.-financial-asset
	possession-owner_2	notHasInPossession	poss.-financial-asset
during situation	possession-theme	hasValue	possession-value

Inherited assertions from ChangeOfPossession:

pre situation	possession-owner_1	hasInPossession	possession-theme
	possession-owner_2	notHasInPossession	possession-theme
post situation	possession-owner_1	notHasInPossession	possession-theme
	possession-owner_2	hasInPossession	possession-theme

EXAMPLES:
 "Marie bought the car from John for 600 dollars"

pre situation	Marie	hasInPossession	600 dollar
	Marie	notHasInPossession	the car
	John	hasInPossession	the car
	John	notHasInPossession	600 dollar
post situation	Marie	hasInPossession	the car
	Marie	notHasInPossession	600 dollar
	John	hasInPossession	600 dollar
	John	notHasInPossession	the car
during situation	the car	hasValue	600 dollar

Figure 3: Non-formal transcription of the mappings, assertions and instantiation for the ESO class FinancialTransaction

from ESO roles (65) to FrameNet Frame Elements (131). The properties in this table pertain to those properties that are used in the situation rule assertions.

3 Predicate Matrix

The PredicateMatrix (PM)(de Lacalle et al., 2014a; de Lacalle et al., 2014b) is an automatic extension of SemLink (Palmer, 2009) that merges several models of predicates such as VerbNet (Kipper et al., 2000), FrameNet (Baker et al., 1998), PropBank (Palmer et al., 2005) and WordNet (Fellbaum, 1998). The PM also contains for each predicate features of the ontologies integrated in the Multilingual Central Repository (Gonzalez-Agirre et al., 2012) like SUMO (Niles and Pease, 2001), Top Ontology (Álvarez et al., 2008) or WordNet domains (Bentivogli et al., 2004). The mappings between such knowledge bases allow to take advantage from their individual strengths. For example, the coverage of PropBank or the semantic relations among events and participants of FrameNet.

The semantic interoperability offered by the PM allows to translate the output of a SRL analysis to

Component	Number
Event classes	63
– DynamicEvent classes	50
– StaticEvent classes	13
SUMO class mappings	46
FrameNet Frame mappings	103
Situation rules	50
Situation rule assertions	123
– Pre situation rule assertions	41
– Post situation rule assertions	52
During situation rule assertions	30
Properties	58
– Unary properties	11
– Binary properties	47
ESO roles	65
Mappings to FrameNet FEs	131

Table 1: Overview of the content in ESO Vers. 2

a representation based on any resource connected to the PM like FrameNet, SUMO or the Domain Ontology. For this reason, we have connected the classes and roles of ESO to the predicates and roles of the PM. We have performed this alignment in two different steps. First, defining a set of new manual mappings between ESO and WordNet. Second, applying an automatic strategy that makes use of the existing mappings between ESO and FN and SUMO. Table 2 contains the number of predicates and roles mapped to ESO by each method.

	Manual	Automatic	Total
predicates	1,702	2,228	3,930
roles	4,831	6,026	10,857

Table 2: Number of predicates and roles mapped to ESO in the PM.

3.1 Manual mappings

For connecting ESO and the PM, manual mappings to Princeton Wordnet 3.0 have been created for all lexical units in a FrameNet frame associated to ESO. In total, 1,614 lexical units from FrameNet have been mapped to WordNet, covering 1,918 synsets. The mappings have been kept outside ESO in order not to overburden the ontology. Additionally, to increase the coverage of ESO in the Predicate Matrix, we manually mapped ESO classes to WordNet Base Level Concepts (BLC). BLCs are important WordNet concepts that cover all WordNet nominal and verbal concepts. In WordNet there are 616 verbal BLCs that cover all 13,151 verbal synsets. The PM can be mapped to 398 of these BLCs which covers 12,722 verbal synsets. The full set of BLCs have been manually checked for their correspondence to an ESO class

and for 75 BLCs a mapping to an ESO class could be made which covers 4,306 synsets.⁵

3.2 Automatic mappings

Both FrameNet and SUMO labels integrated in ESO are used to connect ESO to the PM. For example, the predicate **sell.01** of PropBank belongs, according to its mappings in the PM, to the frame *Commerce_sell* of FrameNet. Consequently, this predicate and its arguments could also be mapped to ESO as shows table 3. Moreover, the frame can also be linked through the SUMO classes. For instance, the predicate **drain.01** of PropBank belongs to the frame *Emptying* that is not considered in ESO. However, it also belongs to the class *Removing* of SUMO and, as a consequence, the mappings in table 4 can be obtained.

PB-pred	PB-arg	FN-frame	FN-fe	ESO-class	ESO-role
sell.01	arg ₀	Commerce_sell	Seller	Selling	possession-owner.1
sell.01	arg ₁	Commerce_sell	Goods	Selling	possession-theme
sell.01	arg ₂	Commerce_sell	Buyer	Selling	possession-owner.2

Table 3: Mapping between PropBank and ESO through FN.

PB-pred	PB-arg	SUMO-class	FN-fe	ESO-class	ESO-role
drain.01	arg ₀	Removing	Theme	Removing	translocation-theme
drain.01	arg ₁	Removing	Source	Removing	translocation-source

Table 4: Mapping between PropBank and ESO through SUMO.

4 Current Output

At the time of submission, about 2.1 million articles on the automotive industry were processed with the NewsReader English pipeline (Agerri et al., 2015) that incorporates the PM and ESO for semantic parsing. Table 5 provides an overview of the number of roles and predicates found, and the number of labels assigned to them per resource in the Predicate Matrix. Note that predicates and roles can receive multiple labels from one resource.

5 Experiment on the WikiNews Corpus

The WikiNews Corpus consists of 120 manually annotated news articles selected from WikiNews⁶ and is used within NewsReader as an evaluation corpus.⁷ The evaluation of the Mate tool that is

⁵All mappings can be downloaded from <https://github.com/newsreader/eso>.

⁶<https://en.wikinews.org>

⁷The corpus will be made available soon at <http://www.newsreader-project.eu/results/data/>

Resource	label frequency
Total predicates	138,695,190
WordNet	293,249,984
VerbNet	236,497,891
PropBank	197,331,322
FrameNet	232,685,360
ESO	85,831,344
Total roles	300,544,817
VerbNet	277,233,904
PropBank	202,134,061
FrameNet	336,248,141
ESO	55,787,300

Table 5: Overview of the number of predicates and roles in a subset of the automotive industry corpus labeled by the Predicate Matrix and ESO

used for the Semantic role labeling scores an F1 of 34.74 for this corpus.⁸ WikiNews has not yet been annotated with ESO classes and roles, as such we used this corpus to test the expressivity and coverage of the Predicate Matrix and ESO first. In short, we followed the same procedure that is also used for the Automotive Corpus. First, all 120 WikiNews articles were processed by the News-Reader Pipeline (Agerri et al., 2015) using the Predicate Matrix and ESO; next, a module called NAF2SEM merged identical events across documents and translated all events into SEM-RDF and finally, all events were loaded into the KnowledgeStore (Corcoglioniti et al., 2015a) and further enriched by the ESO reasoner that infers all ESO assertions, based on the class and role labels. In the KnowledgeStore, the data can be queried via SPARQL queries or simple look-ups.

In table 6 we provide an overview of the results of the first step, the output of the pipeline with respect to the labels for roles and events found. In total, 7,060 predicates were found in the WikiNews corpus. These predicates are assigned one or multiple labels by the Predicate Matrix such as WordNet synset IDs (15,157), FrameNet Frames (12,330) and ESO classes (3,405). The relatively low number of predicates with an ESO class is due to the fact that ESO covers a limited set of concepts and ignores e.g. all speech acts. This table also shows the number of labels found for the roles. In total, 15,652 roles were found that each can again have one or multiple labels.

Next, we derived some basic statistics from the KnowledgeStore that contains all events derived from the corpus. In table 7 we provide an

⁸see (Agerri et al., 2015) for an overview and discussion of these results

Resource	Label frequency
Total predicates	7,060
WordNet	15,157
VerbNet	12,294
PropBank	10,018
FrameNet	12,330
ESO	4,337
Total roles	15,652
VerbNet	14,474
PropBank	10,312
FrameNet	17,680
ESO	3,230

Table 6: Overview of the number of predicate and role labels in the WikiNews corpus labeled by the Predicate Matrix enriched with ESO

Component	Number
Events	5443
ESO events	2508
ESO events with ESO roles	736
ESO events with pre and post situations	444
ESO events with at least one inferred situation	498
ESO events with a during situation	52

Table 7: ESO related statistics of the populated KnowledgeStore of the WikiNews corpus

overview. As is shown, 5,443 distinct events were found of which 2,508 events with an ESO class. Of these events, 736 have at least also an ESO role which is necessary to trigger the situation rules defined in ESO. In total, 444 events were found with inferred pre and post situations and 52 events with inferred during situations. Note that the number of ESO classes that trigger a during situation is smaller (12) than the set of classes that can trigger pre and post situations (46).

Finally, we manually inspected 52 ESO events in the KnowledgeStore with both a pre and post situation (43) and ESO events with a during situation (9).⁹ For this, we randomly selected one or two ESO events per class, depending on the number of occurrences. The result of this inspection are shown in table 8. We found 37 events (71.1%) with a correct class label and 18 events (41.8%) with correct pre and post situations, meaning that the assertions made sense with respect to the original sentences in the document and that the correct role instances were found, if applicable. The set of events with a during situation was correct in 66,6% of the cases. Overall, 21 out of 52 inspected ESO events were found to be correct.

Additionally, we performed an error analysis

⁹The data and analysis can be found at <https://github.com/newsreader/eso>

ESO events with pre/post or during situation	495
Number of events inspected	52 (10.5%)
Number events insp. with a pre/post situation	43
Number events insp. with a during situation	9
Correct class label	37 (71.1%)
Correct pre and post situation(s)	18 (41.8%)
Correct during situation(s)	6 (66.6%)
Correct ESO events	21 (50%)

Table 8: Results of the analysis of ESO events with during or pre/post situation assertions derived from the WikiNews corpus

Error in interpretation sentence (multiple causes)	3
Error in interpretation predicate	9
Multiple conflicting ESO classes assigned	8
Wrong role instance (non-entities)	5
Wrong role instance (entities)	10
Role instance duplication	6
Conflicting assertions	1

Table 9: Results of the error analysis of the inspected ESO events derived from the WikiNews corpus

to investigate where errors or omissions stemmed from. The results of the error analysis can be found in table 9. In general, each of the 16 modules in the pipeline introduces some errors, which is reflected in the outcome of the error analysis. For nine events we found that the sense of the predicate was misinterpreted, for eight events multiple and conflicting ESO classes were assigned due to some unavoidable level of ambiguity in the Predicate Matrix. In five cases, we found that the Semantic Role Labeler picked up the wrong role; for ten events DBpedia Spotlight assigned a wrong label for a named entity. These errors also resulted in 6 role duplications where subject and object of an assertion are identical while they should not. For one event, it caused conflicting assertions.

6 Discussion and Future Work

In this paper, we have presented the new release of the Event and Implied Situation Ontology (ESO) and the PredicateMatrix (PM). Both resources augment Semantic Role Labeling techniques and are applied to a very large document collection to capture implications of events for a selected set of concepts, roles and properties. Through the WordNet backbone of ESO and PM, we were able to derive a formal model for event implications with a large coverage in English. Since wordnets in many languages are connected to WordNet, this model has also been projected to other languages in the

NewsReader project: Spanish, Dutch and Bulgarian. ESO thus has shown to be used as a interoperable framework on reasoning over changes and their implications across different languages. This allows us to compare the content of text across languages, regardless of the way this content is expressed.

From the experiment on the WikiNews corpus, we conclude that ESO performs reasonably well on this dataset with 50% of correct ESO events with a pre/post or during situation. The ontology is not built in order to define all events in text which is shown in the coverage of all events found (5,443), and the ESO events (2,508) of which 496 have either both a pre and post situation or a during situation. The errors in the ESO events with assertions are mainly caused by an unavoidable degree of errors in the processing pipeline as was reported in the error analysis. The observation that not all ESO events come with assertions is likely due to the fact that a sentence does not always contain all roles necessary for an assertion rule to fire. A more in-depth analysis of the annotated texts will provide an answer for this.

We are currently processing about 2.1 million news articles on the automotive industry, where the ESO mapping are inserted in the SRL layers. The output is converted to RDF, after which we apply reasoning to derive new statements as was shown in the experiment. The output will be evaluated through inspecting samples, against benchmark data that will be developed on the WikiNews corpus and through end-user tasks on the data sets. Also, we planned additional experiments on the usability of the ESO assertions for tracking actual chains of property changes through time. Finally, the WikiNews corpus has been translated to Spanish, Dutch and Italian. The processing of the translated text through NewsReader pipelines in these languages, where these pipelines exploit the same ESO model and a language-specific PM, will allow us to do a cross-language comparison of the inferred properties.

Acknowledgments

The research for this paper was supported by the European Union’s 7th Framework Programme via the NewsReader Project (ICT-316404).

References

- V. Presutti, A. Nuzzolese, A. Gangemi. 2012. Gathering lexical linked data and knowledge patterns from FrameNet. In *Proceedings of K-CAP '11*.
- R. Agerri, I. Aldabe, Z. Beloki, E. Laparra, M. Lopez de Lacalle, G. Rigau, A. Soroa, A. Fokkens, R. Izquierdo, M. van Erp, P. Vossen, C. Girardi, and A. Minard. 2015. Event detection, version 3. Deliverable 4.2.3. NewsReader-ICT316404.
- J. Álvarez, J. Atserias, J. Carrera, S. Climent, A. Oliver, and G. Rigau. 2008. Consistent annotation of EuroWordNet with the Top Concept Ontology. In *Proceedings of GWC'08*.
- C.F. Baker, C.J. Fillmore, and J.B. Lowe. 1998. The Berkeley FrameNet project. In *Proceedings COLING-ACL, ACL '98*, Montreal, Canada.
- L. Bentivogli, P. Forner, B. Magnini, and E. Pianta. 2004. Revising the Wordnet domains hierarchy: semantics, coverage and balancing. In *Proceedings of the Workshop on Multilingual Linguistic Resources*. ACL.
- F. Corcoglioniti, M. Rospocher, R. Cattoni, B. Magnini, and L. Serafini. 2015a. The KnowledgeStore: a storage framework for interlinking unstructured and structured knowledge. *International Journal on Semantic Web and Information Systems*, 11(2):1–35, April-June.
- F. Corcoglioniti, M. Rospocher, M. Mostarda, and M. Amadori. 2015b. Processing billions of RDF triples on a single machine using streaming and sorting. In *ACM SAC 2015 Proceedings*.
- M. López de Lacalle, E. Laparra, and G. Rigau. 2014a. First steps towards a Predicate Matrix. In *Proceedings of GWC 2014*, Tartu, Estonia.
- M. López de Lacalle, E. Laparra, and G. Rigau. 2014b. Predicate matrix: extending SemLink through WordNet mappings. In *Proceedings of LREC'14*, Reykjavik, Iceland.
- C. Fellbaum. 1998. *WordNet: an electronic lexical database*. MIT Press.
- A. Gonzalez-Agirre, E. Laparra, and G. Rigau. 2012. Multilingual central repository version 3.0. In *Proceedings of LREC '11*, pages 2525–2529.
- A. Hicks. 2009. Domain extension of the central ontology - final. Deliverable 8.3, KYOTO-ICT 211423.
- S. Im and J. Pustejovsky. 2009. Annotating event implicatures for textual inference tasks. In *Proceedings of the 5th International Conference on Generative Approaches to the Lexicon*.
- S. Im and J. Pustejovsky. 2010. Annotating lexically entailed subevents for textual inference tasks. In *Proceedings of FLAIRS-23*, Daytona Beach, USA.
- K. Kipper, H. Trang Dang, and M. Palmer. 2000. Class-based construction of a verb lexicon. In *Seventeenth National Conference on Artificial Intelligence, AAAI-2000*.
- K. Kipper, A. Korhonen, N. Ryant, and M. Palmer. 2006. Extending VerbNet with novel verb classes. In *Proceedings of LREC 2006*.
- C. Masolo, S. Borgo, A. Gangemi, N. Guarino, A. Oltramari, and L. Schneider. 2002. Wonderweb deliverable d17. Technical report, ISTC-CNR.
- I. Niles and A. Pease. 2001. Towards a standard upper ontology. In *Proceedings of FOIS-Volume 2001*. ACM.
- M. Palmer, D. Gildea, and P. Kingsbury. 2005. The proposition bank: An annotated corpus of semantic roles. *Computational Linguistics*, 31(1):71–106.
- M. Palmer. 2009. Semlink: Linking Propbank, Verbnet and Framenet. In *Proceedings of the Generative Lexicon Conference*, pages 9–15.
- J. Scheffczyk, A. Pease, and M. Ellsworth. 2006. Linking FrameNet to the Suggested Upper Merged Ontology. In *Proceedings of FOIS 2006*.
- R. Segers, P. Vossen, M. Rospocher, L. Serafini, E. Laparra, and G. Rigau. 2015. ESO: a frame based ontology for events and implied situations. In *Proceedings of Maplex 2015*, Yamagata, Japan.