The Event and Implied Situation Ontology (ESO): Application and Evaluation

Roxane Segers*, **Marco Rospocher** ∓, **Piek Vossen** *, **Egoitz Laparra** †, **German Rigau** †, **Anne-Lyse Minard**∓

*Vrije Universiteit Amsterdam, The Netherlands
{r.h.segers, p.t.j.m.vossen}@vu.nl
∓Fondazione Bruno Kessler, Trento, Italy
{rospocher, minard,}@fbk.eu
† University of the Basque Country (UPV/EHU), Spain
{egoitz.laparra, german.rigau}@ehu.eus

Abstract

This paper presents the Event and Implied Situation Ontology (ESO), a manually constructed resource which formalizes the pre and post situations of events and the roles of the entities affected by an event. The ontology is built on top of existing resources such as WordNet, SUMO and FrameNet. The ontology is injected to the Predicate Matrix, a resource that integrates predicate and role information from amongst others FrameNet, VerbNet, PropBank, NomBank and WordNet. We illustrate how these resources are used on large document collections to detect information that otherwise would have remained implicit. The ontology is evaluated on two aspects: recall and precision based on a manually annotated corpus and secondly, on the quality of the knowledge inferred by the situation assertions in the ontology. Evaluation results on the quality of the system show that 50% of the events typed and enriched with ESO assertions are correct.

Keywords: Ontology, Semantic Role Labeling, Text Mining, Semantic Web

1. Introduction

In this paper, we present the Event and Implied Situation Ontology (ESO) Version 2, that is matched with 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, thus favoring the construction of large event-centric knowledge graphs (Rospocher et al., to appear).

ESO Version 2 (Segers et al., 2016) is a newly developed domain ontology to enhance the extraction and linking of dynamic and static events and their implications in text. Such a chain of changes and states and their implied situations is presented in Figure 1. Here, the boxes represent various event expressions about John's employment while the ovals represent the implied situations of each event. Modeling of event implications allows for extracting sequences of states and changes over time regardless of this information being directly expressed in text, or inferred by a reasoner. The model targets interpretations of situations rather than the semantics of predicates per se. Events are interpreted as situations 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.

Lexicons that define implications of events, for example VerbNet (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 the event components to make these implications explicit. Though some research on deductive reasoning over Frame annotated text (Scheffczyk et al., 2006) and defining pre and post situations of predicates exist (Im and Pustejovsky, 2009), to the best of our knowledge, ontologies that model events, roles and implications for semantic parsing do not. 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. Also, SUMO needs a sophisticated reasoning system to be productive (Álvez et al., 2015). Furthermore, such ontologies need to be integrated with semantic parsing systems that deal with expressions on natural language. We therefore decided to develop a new ontology for modelling events and their implications that is tailored to a semantic parsing system for text.

The Predicate Matrix (López de Lacalle et al., 2014) integrates predicate and role information from FrameNet (Baker et al., 1998), VerbNet (Kipper et al., 2000), PropBank (Palmer et al., 2005), NomBank (Meyers et al., 2004) and WordNet (Fellbaum, 1998). This resource is used to assign role and predicate annotations at sentence level. All classes and roles in ESO are fed back into the Predicate Matrix. As such the ontology provides an additional layer of annotations in text that allow for inferencing over events and implications.

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 an overview of the Predicate Matrix and ESO in our document collection. In Section 5. we report on the evaluation of the ontology. We conclude in Section 6. with a discussion and some outlines for future work.

2. ESO: Meta Model and Content

In this paragraph, we briefly describe the meta model of the ontology and provide an overview of its content. The ESO

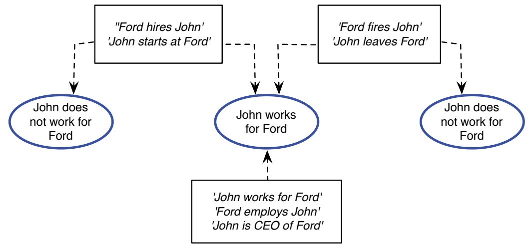

Figure 1: A chain of expressions of changes and states and their implied situations

ontology and documentation can be found online: `https://github.com/newsreader/eso`

ESO is an OWL 2 ontology.[1] 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:JoiningAnOrganization;
        eso:employment-employee    :John;
        eso:employment-employer    :Ford;
        eso:employment-function   :new CEO;
        sem:hasTime  :time_eventX.
}
```

specifies that the event (X) is of a certain type (eso:JoiningAnOrganization), that it involves an entity playing the role of employee (:John), an entity playing the role of employer (:Ford), an entity playing the role of function (:new CEO) 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, five core classes are defined in ESO: 1) **Event**: this class is the root of the taxonomy of event types. Any event detected in a text is 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 are instantiated starting from the event instances detected in the text; 5) **SituationRule**: the individuals of this class 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 ontology. This is then exploited by a reasoning module that instantiates 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:JoiningAnOrganization has two specific individuals: pre_JoiningAnOrganization and post_JoiningAnOrganization. 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 eso:JoiningAnOrganization, the individual pre_JoiningAnOrganization has one eso:SituationRuleAssertion, where eso:pre_JoiningAnOrganization_assertion_1 states:

```
eso:pre_JoiningAnOrganization_assertion1
    eso:hasSituationAssertionSubject      eso:employment-employee;
    eso:hasSituationAssertionProperty     eso:notEmployedAt;
    eso:hasSituationAssertionObject       eso:employment-employer.
```

Based on the class assertions, the ESO reasoner[2] can now infer that the event belongs to the class eso:JoiningAnOrganization and that it has entity instances in certain roles where some entity does not work for some employer before the event and that he works for some employer and in some function after the event. In the case of the example sentence 'Ford hires John as their new CEO', the instantiation of the defined situations for the event instance of eso:JoiningAnOrganization will then look as follows:

```
:eventX_pre {
     :John        eso:notEmployedAt      :Ford
:eventX_post
     :John        eso:employedAt         :Ford
     :John        eso:hasFunction     :new CEO
     :John        eso:isEmployed         : 'true'
}
```

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 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:Decreasing, 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:Decreasing, the existential restriction is defined as follows:

```
eso:Decreasing  rdfs:subClassOf [
    a owl:Restriction ;
    owl:onProperty        eso:triggersPreSituationRule    ;
    owl:hasValue          eso:pre_Decreasing ] .
eso:Decreasing  rdfs:subClassOf [
    a owl:Restriction ;
    owl:onProperty        eso:triggersPostSituationRule   ;
    owl:hasValue          eso:post_Decreasing ] .
eso:Decreasing  rdfs:subclassOf [
    a owl:Restriction ;
    owl:onProperty       eso:quantity-attribute    ;
    owl:someValuesFrom   owl:Thing] .

eso:pre_Decreasing    a eso:SituationRule .
eso:post_IDecreasing   a eso:SituationRule .
```

[1] `http://www.w3.org/2001/sw/wiki/OWL`

[2] Implemented as a processor of RDFpro (Corcoglioniti et al., 2015b). See also: `http://bit.ly/ESOreasoner`

These are the situation rule assertions defined for the pre an post situation of eso:Decreasing:

```
eso:pre_Decreasing_assertion1
    eso:hasSituationAssertionSubject      eso:quantity-item;
    eso:hasSituationAssertionProperty     eso:hasAttribute;
    eso:hasSituationAssertionObject       eso:quantity-attribute.

eso:pre_Decreasing_assertion2
    eso:hasSituationAssertionSubject      eso:quantity-attribute;
    eso:hasSituationAssertionProperty     eso:hasRelativeValue;
    eso:hasSituationAssertionObjectValue  `+'

eso:post_Decreasing_assertion1
    eso:hasSituationAssertionSubject      eso:quantity-item;
    eso:hasSituationAssertionProperty     eso:hasAttribute;
    eso:hasSituationAssertionObject       eso:quantity-attribute.

eso:post_Decreasing_assertion2
    eso:hasSituationAssertionSubject      eso:quantity-attribute;
    eso:hasSituationAssertionProperty     eso:hasRelativeValue;
    eso:hasSituationAssertionObjectValue  `-'
```

The pre and post situation named graphs for the example sentence ”Ford decreased 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 less (-) after the event than it was before the event (+), meaning that the production goes from more (+) to less (-).

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

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

and the assertions will be instantiated as:

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

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. In many cases, textual description are underspecified with respect to the specific attrbute that is implied. Requiring these attributes to be present means we cannot use the information at all. 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. Content of ESO

ESO is a hand-built resource, based on high-frequent FrameNet frames that were extracted from a large domain-specific document collection. We mapped these frames manually to the SUMO ontology [3], to derive an initial conceptual structure. 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.

[3] http://www.ontologyportal.org

```
-FinancialTransaction:  subclassOf: ChangeOfPossession
    "The subclass ofChangeOfPossession 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 2: Informal transcription of the mappings, assertions and instantiation for the ESO class FinancialTransaction

Besides classes, the other main components of the ontology consist 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. As such, not all FEs of a frame are mapped to ESO but only those that play a role in modeling the situation assertions.

To illustrate the expressivity of the situation assertions, in Figure 2 we provide an informal transcription of a typical class in ESO, including the class mappings to SUMO and FrameNet, the aggregated role mappings to FrameNet FEs, the inherited and class specific situation assertions and an example of the instantiation. From the knowledge in the example sentence ”Marie bought the car from John for 600 dollars”, we are able to infer that a) Marie has 600 dollar and not this car *before* the event, while John does have this car but not the 600 dollar; b) *after* the event, the money and the car have changed of ownership while c) the car itself has a value of 600 dollar *during* the exchange.

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

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 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. The Predicate Matrix

The PredicateMatrix (PM) [4] is an automatic extension of SemLink (Palmer, 2009) that merges several models of predicates such as VerbNet, FrameNet, PropBank and WordNet. 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 (Álvez et al., 2008) or WordNet domains (Bentivogli et al., 2004).
The semantic interoperability offered by the PM allows to translate the output of a SRL analysis to 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 manual mappings between ESO and WordNet. Second, applying an automatic strategy that makes use of the existing mappings between ESO and FrameNet 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.

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 different ESO classes. It also contains 23,386 roles of PropBank and NomBank mapped to 2,343 frame-elements of FrameNet and 53 ESO roles.

[4] http://adimen.si.ehu.es/web/PredicateMatrix

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

4. Current output of the system

About 2.3 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 3 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. Evaluation

The ESO ontology and the Predicate Matrix were evaluated on three aspects. First, we derived the recall and precision of the ESO labeling of predicates and roles by the NewsReader pipeline compared to a manually annotated corpus and a baseline system. Next, we focus on the *quality* of the events and the knowledge inferred by the assertion rules in ESO for the NewsReader version of the corpus, a Gold Standard and a baseline system. Finally, we evaluated a sample of ESO events in the 2.3 million Automotive Corpus.

5.1. Manual annotation of articles with ESO classes and roles

The MEANTIME Corpus (Minard et al., 2016) consists of 120 news articles selected from WikiNews [5] and is used as an evaluation corpus for the NewsReader NLP pipeline.[6] To evaluate the labeling with ESO via the Predicate Matrix, this corpus was annotated with ESO classes and roles. If a predicate has no participants, we only typed the predicate with an ESO class. Annotations were applied for the first five sentences of each article, thus resulting in an annotated corpus of 600 annotated sentences.[7] The corpus was annotated by one person; on a subset of twelve articles annotated by two persons we derived an Inter Annotator Agreement of 0.65 for the predicates and 0.60 for the roles.

[5] https://en.wikinews.org

[6] The corpus is available at http://www.newsreader-project.eu/results/data/wikinews

[7] The MEANTIME corpus with ESO annotations can be found at https://github.com/newsreader/eso

	Predicates	Roles
Precision	36.1%	28.2%
Recall	68.2%	52.8%

Table 4: Results of the NewsReader system versus the Gold Standard

	Predicates	Roles
Precision	61.6%	34.3%
Recall	37.5%	27.2%

Table 5: Results of the baseline system versus the Gold Standard

NewsReader system versus Gold Standard In total, 712 predicates in MEANTIME were manually annotated with an ESO class and 1033 roles (3024 tokens) with an ESO role. The system labeled 5655 role tokens and 1344 predicates. In total, 486 predicates have been labeled with the correct ESO class and 1597 tokens with the correct ESO role. The evaluation results are presented in Table 4

The precision and recall of the predicates is 36.1% and 68.2% respectively; for the roles, the precision and recall is 28.2% and 52.8%. The evaluation of the Mate tool that is used for the Semantic role labeling scores 20.9 on precision and 80.2 on recall for this corpus.[8] These results are partially due to the fact that ESO was designed for the automotive industry domain, while MEANTIME comprises news about other industries such as aviation and stock markets. The low precision is explained by the fact that the SRL module and the Predicate Matrix tend to overgenerate for example by assigning multiple ESO classes to one predicate or by assigning the same role to different entity mentions. Also, the SRL tends to recognize NPs such as 'market' as a predicate and tries to find roles for this predicate where there are none. The low recall is due to the way the corpus was annotated. For the annotation we also included predicates and roles of which we knew that they were very difficult to find by the system, for example because of nested clauses or metaphorical language use. Additionally, not all roles defined in ESO have a mapping in the Predicate Matrix.

Baseline system versus Gold Standard As a baseline system, we stripped the Predicate Matrix down to the original SemLink (Palmer, 2009) elements and preserved the remaining mappings to ESO. The results are shown in Table 5. SemLink plus ESO labeled a total of 2405 role tokens and 433 predicates. For 267 predicates the ESO class is correct and 825 tokens are labeled with the correct ESO role.

The precision and recall of the predicates is 61.6% and 34.3% respectively; for the roles the precision and recall is 34.3% and 27.2%. The Predicate Matrix plus ESO outperforms the baseline system on recall, while the baseline system outperforms on precision. As was mentioned, the SRL and Predicate Matrix tend to overgenerate which is reflected in the number of predicates labeled with ESO by SemLink (433) and the Predicate Matrix (1344) while 712 predicates were labeled with and ESO class in the Gold Standard.

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

Resource	Label frequency KS-1
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: KS-1: Overview of the number of predicate and role labels in the MEANTIME corpus labeled automatically with the Predicate Matrix enriched with ESO

5.2. Quality checks of the automatically derived ESO events and assertions

For the quality checks of the ESO events and situations, we created three KnowledgeStores[9] (Corcoglioniti et al., 2015a); the first one (KS-1) is based on the output of the Newsreader pipeline on the MEANTIME corpus, the second one (KS-2) is based on the Gold Standard manual annotations of MEANTIME. The third KnowledgeStore (KS-3) is based on the output of the baseline system (Semlink plus ESO).

For the automatic version of the KnowledgeStore (KS-1), all 120 articles in the MEANTIME corpus were processed by the NewsReader 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. For the Gold Standard version of the KnowledgeStore (KS-2), we took the manual annotations as input, converted the annotated documents to NAF, ran NAF2SEM and loaded the events in the KnowledgeStore. The same procedure was followed for the baseline system (KS-3).

5.2.1. Quality checks of the NewsReader system in KS-1

In Table 6 we provide the results of the output of the NewsReader pipeline with respect to the labels for roles and events found. In total, 7,060 predicates were found in the MEANTIME 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 (4,337). 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 all speech acts. This table also shows the number of labels found for the roles. In total, 15,652 roles were found.

[9]`https://knowledgestore.fbk.eu/`

Component	KS-1
Events	5,443
ESO events	2,508
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: KS-1: ESO related statistics of the populated KnowledgeStore

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 MEANTIME corpus

Next, we derived basic statistics from the KS-1 that contains all *events* derived from the corpus. In Table 7 we provide an 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).[10] For this, we randomly selected one or two ESO events per class, depending on the number of occurrences. The results 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 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 SRL 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.

[10]The data and analysis can be found at `https://github.com/newsreader/eso`

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 MEANTIME corpus

Resource	Label frequency KS-1	Label frequency KS-2
ESO predicates	4,337	712
ESO roles	3,230	1033

Table 10: KS-2: Overview of the number of ESO predicates and roles in the manually annotated MEANTIME corpus

5.2.2. Quality checks of the GoldStandard in KS-2

KnowledgeStore-2 is based on the manual annotations of MEANTIME with ESO classes and roles. For this corpus, we provide the number of predicate and role labels in Table 10. In this table we also provide the numbers pertaining to KS-1 for comparison. As is shown, the amount of manual annotated ESO predicates and roles is far less than those generated by the system. As was explained, this is mainly due to the fact that the SRL assigns multiple ESO classes for a predicate. Also, the system labels words as predicate (e.g. "manufacturer") which are not predicates. For these cases, roles are generated which is reflected in higher number of role labels in KS-1 than in KS-2. Note that MEANTIME was not annotated with WordNet, PropBank, FrameNet or VerbNet.

Next, we derived basic statistics from the KS-2; in Table 11 we provide the numbers of KS-1 and KS-2 for comparison. As is shown, 1,120 distinct events were found of which 441 events with an ESO class. Of these events, 406 are accompanied with at least one ESO role. In total, 268 events were found with inferred pre and post situations and 47 events with inferred during situations. These results show that less events were found based on the manually annotated data than in the automatically generated data.

The ESO events in KS-2 were inspected manually as was done for KS-1 to check a) the quality of the assertions and

Component	KS-1	KS-2
Events	5,443	1,120
ESO events	2,508	441
ESO events with ESO roles	736	406
ESO events with pre and post situations	444	268
ESO events with at least one inferred situation	498	320
ESO events with a during situation	52	47

Table 11: ESO related statistics of the populated KnowledgeStores (KS-1 and KS-2) based on the MEANTIME corpus

Component	KS-1	KS-2	KS-3
Events	5443	1120	5557
ESO events	2508	441	993
ESO events with ESO roles	736	406	622
ESO events with pre and post situations	444	268	342
ESO events with at least one inferred situation	498	320	361
ESO events with a during situation	52	47	19

Table 12: ESO related statistics of the baseline KnowledgeStores (KS-3)

b) to check whether any errors were introduced by either the NAF2SEM module or the reasoner. As such we inspected 25 events with pre and post situations and 5 events with a during situation. Since these ESO events are based on manual annotations, the quality of the events, assertions and roles reaches 92%. The three minor errors that were found were caused by the event coreference module.

5.2.3. Quality checks of the baseline system in KS-3

KnowledgeStore-3 is based on SemLink plus ESO annotations of MEANTIME. For this corpus, we present the basic statistics in Table 12; we also provide the numbers of KS-1 and KS-2 for comparison.

As is shown in these results, the baseline system finds about as many events (5557) as the SRL combined with the Predicate Matrix, but generates far less events labeled with an ESO class (993). For this baseline KnowledgeStore, we inspected 25 events with a pre and post situation and 5 events with a during situation. An event was typed with the correct ESO label in 63% of the cases; 36% of the events have the correct class, roles and inferred situations. As such, the ESO events and assertions in the KnowledgeStore based on the Predicate Matrix (KS-1) outperform the baseline system on both correct typing (71.1% in KS-1) and correct typing, roles and assertions (50% in KS-1).

5.3. Stratified quality checks on the Automotive corpus

Next to the evaluation of ESO and the Predicate Matrix on the MEANTIME corpus, we carried out additional quality checks on a sample of 100 events taken from the Automotive corpus. We based the sample on the four different heuristics used to map ESO to the Predicate Matrix: the set ESO1 is based on the manual mappings to Princeton WordNet; ESO2 is based on the manual Base Level Concept mappings; the set ESO3 is based on automatic mappings via SUMO and finally, the set ESO4 is based on automatic FrameNet mappings.

For each ESO event the class was checked as well as the correctness of the roles and assertions, if any. The results are presented in Table 13. The table shows that on average 61% of the ESO events have the correct class and in 39% the ESO events have the correct assertions and roles. Further, 28% of the ESO events in the samples come without any assertions due to the fact that not all necessary roles

Set	ESO class correct	ESO class incorrect	ESO events with correct assertions	ESO events without assertions
ESO1	15 (60%)	10 (40%)	12 (48%)	7 (28%)
ESO2	18 (72%)	7 (28%)	8 (32%)	7 (28%)
ESO3	13 (52%)	12 (48%)	9 (36%)	5 (20%)
ESO4	15 (60%)	10 (40%)	7 (28%)	9 (36%)
Average	61%	39%	36%	28%

Table 13: Results of the quality checks of ESO events in the Automotive Corpus

were found in the text to fire the assertions rules. If we compare these result to the quality checks carried out on the MEANTIME corpus, we see about a 10% drop in the percentage of correct classifications and correct full ESO events. This is possibly due to the fact that we deliberately included ESO events in this sample without any assertions; some of the noise generated by the SRL module is cleared out if an ESO event comes with roles and assertions.

6. Discussion and Future Work

We presented ESO and its integration into the Predicate Matrix; both resources augment Semantic Role Labeling techniques. They are applied to a very large document collection to capture implications of events for a selection of concepts, roles and properties. Both resources have also been projected to other languages: Spanish, Dutch and Bulgarian (López de Lacalle et al., 2016). As such, ESO can 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.

The evaluation results on MEANTIME and the Automotive corpus show that ESO events and their asserted situations show promising results with respect to quality, even though the NewsReader pipeline yields a relatively low precision and recall for ESO related predicates and roles.

Future work will include the evaluation of the *relevance* of the asserted situations. For this we planned to compare manually and automatically generated event timelines derived from the MEANTIME corpus. Also, we planned additional experiments on the usability of the ESO assertions for tracking actual chains of property changes through time in the Automotice corpus. Further development of the ontology itself will encompass mapping ESO to DOLCE and developing additional role restrictions to constrain the selection of instances.

Acknowledgments

The research for this paper has been partially funded by the European Union 7th Framework Programme NewsReader (FP7-ICT-2011-8-316404) and the Spanish national project TUNER (TIN2015-65308-C5-1-R).

7. References

Agerri, R., Aldabe, I., Beloki, Z., Laparra, E., Rigau, G., Soroa, A., van Erp, M., Fokkens, A., Ilievski, F., Izquierdo, R., Morante, R., van Son, C., Vossen, P., and Minard, A.-L. (2015). Event detection, version 3. Deliverable 4.2.3. NewsReader-ICT316404.

Álvez, J., Atserias, J., Carrera, J., Climent, S., Oliver, A., and Rigau, G. (2008). Consistent annotation of EuroWordNet with the Top Concept Ontology. In *Proceedings of GWC'08*.

Álvez, J., Lucio, P., and Rigau, G. (2015). Improving the Competency of First-Order Ontologies. In *Proceedings of K-CAP 2015*.

Baker, C., Fillmore, C., and Lowe, J. (1998). The Berkeley FrameNet Project. In *Proceedings COLING-ACL*, ACL '98, Montreal, Canada.

Bentivogli, L., Forner, P., Magnini, B., and Pianta, E. (2004). Revising the Wordnet Domains Hierarchy: Semantics, Coverage and Balancing. In *Proceedings of the Workshop on Multilingual Linguistic Resources*, MLR '04, pages 101–108, Stroudsburg, PA, USA. Association for Computational Linguistics.

Corcoglioniti, F., Rospocher, M., Cattoni, R., Magnini, B., and Serafini, L. (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.

Corcoglioniti, F., Rospocher, M., Mostarda, M., and Amadori, M. (2015b). Processing Billions of RDF Triples on a Single Machine using Streaming and Sorting. In *Proceedings of the 30th Annual ACM Symposium on Applied Computing*, SAC '15, pages 368–375, New York, NY, USA. ACM.

Fellbaum, C. (1998). *WordNet: an electronic lexical database*. MIT Press.

Gonzalez-Agirre, A., Laparra, E., and Rigau, G. (2012). Multilingual Central Repository version 3.0. In *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC'12)*, Istanbul, Turkey, may. European Language Resources Association (ELRA).

Im, S. and Pustejovsky, J. (2009). Annotating Event Implicatures for Textual Inference Tasks. In *The 5th Conference on Generative Approaches to the Lexicon*.

Kipper, K., Dang, H. T., and Palmer, M. (2000). Class-Based Construction of a Verb Lexicon. In *17th National Conference on Artificial Intelligence*.

Kipper, K., Korhonen, A., Ryant, N., and Palmer, M. (2006). Extending VerbNet with Novel Verb Classes. In *Proceedings of LREC 2006*.

López de Lacalle, M., Laparra, E., and Rigau, G. (2014). Predicate Matrix: extending SemLink through WordNet mappings. In *Proceedings of LREC'14*.

López de Lacalle, M., Laparra, E., Aldabe, I., and Rigau, G. (2016). A Multilingual Predicate Matrix. In *Proceedings of the 10th International Conference on Language Resources and Evaluation (LREC'16)*, Portoroz, Slovenia, May. European Language Resources Association (ELRA).

Masolo, C., Borgo, S., Gangemi, A., Guarino, N., Oltramari, A., and Schneider, L. (2002). WonderWeb Deliverable D17. Technical report, ISTC-CNR.

Meyers, A., Reeves, R., Macleod, C., Szekely, R., Zielinska, V., Young, B., and Grishman, R. (2004). The NomBank project: An interim report. In *HLT-NAACL 2004 workshop: Frontiers in corpus annotation*, pages 24–31.

Minard, A., Speranza, M., Urizar, R., Altuna, B., van Erp, M., Schoen, A., and van Son, C. (2016). MEANTIME, the NewsReader Multilingual Event and Time Corpus. In *Proceedings of LREC2016*.

Niles, I. and Pease, A. (2001). Towards a Standard Upper Ontology. In *Proceedings of FOIS-Volume 2001*. ACM.

Palmer, M., Gildea, D., and Kingsbury, P. (2005). The Proposition Bank: An Annotated Corpus of Semantic Roles. *Computational Linguistics*, 31(1):71–106, March.

Palmer, M. (2009). Semlink: Linking Propbank, VerbNet and FrameNet. In *Proceedings of the Generative Lexicon Conference*, pages 9–15.

Rospocher, M., van Erp, M., Vossen, P., Fokkens, A., Aldabe, I., Rigau, G., Soroa, A., Ploeger, T., and Bogaard, T. (to appear). Building Event-Centric Knowledge Graphs from News. *Journal of Web Semantics*.

Scheffczyk, J., Pease, A., and Ellsworth, M. (2006). Linking FrameNet to the Suggested Upper Merged Ontology. In *Proceedings of FOIS 2006*.

Segers, R., Laparra, E., Rospocher, M., Vossen, P., Rigau, G., and Ilievski, F. (2016). The Predicate Matrix and the Event and Implied Situation Ontology: Making more of Events. In *Proceedings of GWC2016*.