Automatically Extracting Qualia Relations for the Rich Event Ontology

Ghazaleh Kazeminejad[1], Claire Bonial[2], Susan Windisch Brown[1] and **Martha Palmer[1]**
[1]University of Colorado Boulder, [2]U.S. Army Research Lab
{ghazaleh.kazeminejad, susan.brown, martha.palmer}@colorado.edu
claire.n.bonial.civ@mail.mil

Abstract

Commonsense, real-world knowledge about the events that entities or "things in the world" are typically involved in, as well as part-whole relationships, is valuable for allowing computational systems to draw everyday inferences about the world. Here, we focus on automatically extracting information about (1) the events that typically bring about certain entities (origins), (2) the events that are the typical functions of entities, and (3) part-whole relationships in entities. These correspond to the agentive, telic and constitutive qualia central to the Generative Lexicon. We describe our motivations and methods for extracting these qualia relations from the Suggested Upper Merged Ontology (SUMO) and show that human annotators overwhelmingly find the information extracted to be reasonable. Because ontologies provide a way of structuring this information and making it accessible to agents and computational systems generally, efforts are underway to incorporate the extracted information to an ontology hub of Natural Language Processing semantic role labeling resources, the Rich Event Ontology.

1 Introduction

Most adults could tell you that an umbrella is used to keep the rain off of you, that bread is baked, or that dandelions are made up of leaves, stems, and flowers. Commonsense, real-world knowledge about the events that entities or "things in the world" are typically involved in, as well as part-whole relationships, is valuable for allowing computational systems to draw everyday inferences about the world. Here, we describe our approach for automatically extracting this information from the Suggested Upper Merged Ontology (SUMO) (Niles and Pease, 2001a; Pease, 2011). We assume the theoretical framework of the Generative Lexicon (GL) (Pustejovsky, 1991), as we find GL qualia relations useful for considering the typical involvement of a particular entity in events. Thus, we focus on relationships between entities and events captured in a set of qualia associated with an entity. We find that human annotators asked to evaluate the quality and accuracy of the information extracted overwhelmingly find the origins, functions and part-whole relationships proposed to be reasonable.

One of our motivations is to add the extracted qualia relations to an ontology hub of Natural Language Processing (NLP) semantic role labeling (SRL) resources, the Rich Event Ontology (REO). We plan to exploit the qualia relations within REO to make generalizations across participant types for a given event type, thereby facilitating the addition of selectional restrictions, and to enable some shallow inferencing about how entities generally come into existence, how they are used, and what they are made up of. By structuring this information within REO we aim, for example, to allow an agent or computational system to infer that a room (whether detected visually or described in text) with a stove, pans, and a refrigerator is likely a kitchen where food preparation happens. Our approach first requires expansion of REO to include additional "things," or ENDURANTS as they are called in REO, since past efforts have focused on the inclusion of events or PERDURANTS. We again draw information from SUMO to quickly and efficiently add high-quality information on ENDURANTS to REO, and discuss the linking model we have developed to integrate the new classes from SUMO.

Proceedings of the 27th International Conference on Computational Linguistics, pages 2644–2652
Santa Fe, New Mexico, USA, August 20-26, 2018.

2 Theoretical Framework: Generative Lexicon

The theory of GL focuses on the distributed nature of compositionality in natural language (Pustejovsky, 1991), which attributes parts of the semantic load to non-verb constituents, as opposed to approaches where verbs are at the center of the notion of compositionality, (e.g. Levin (1993)). According to GL theory, every entity has a set of qualia that plays a role in what events that entity can be involved in, and that directly affects the compositional meaning that arises from the mention of a particular entity being involved in a particular event. The set of qualia GL proposes include formal, telic, agentive, and constitutive. The formal quale represents the IS-A relation for a given entity. For example, the formal quale for *dandelion* might be *plant*. The constitutive quale is the set of other entities that are typically considered to be parts of the given entity or the material of the entity. For example, the constitutive quale for *plant* could be *leaf*, *stem*, etc., and the constitutive quale for *pottery* could be *clay.* The telic quale represents the typical purpose or function of the given entity, if there is one (e.g., the telic quale of *umbrella* might be *protect from rain*). The agentive quale represents the factors involved in how the given entity came into being (e.g., the agentive quale of *bread* could be *baking*).

Note that there is no single set of defined quale for the entities of the world, hence there is indeterminacy in how these should be labeled and at what level of granularity (i.e., is the agentive quale of *bread* better captured as *baking* or as *applying heat*?). It is in part for this reason that situating qualia relations in an ontology, as was done with SIMPLE (Lenci et al., 2000), is so appealing. An ontology provides a common vocabulary for entities and events and structures them in a hierarchy from most general to most specific.

In an ontology with GL qualia, any node would inherit all the qualia from the set of its ancestors. The deeper we go into the lower levels of the hierarchy, the more specific the qualia get, and, in some cases, they may override parts of the qualia they have inherited. For instance, any physical object has a physical form (formal quale), so any `Artifact` inherits this property because it is a child to the physical object node in the ontology. In addition to this inherited quale, `Artifact` introduces to the ontology its own set of qualia: an agentive quale of *making* and a very general telic quale of *some event,* since every artifact is 'made' and typically has 'some purpose'.

3 Resource Background

3.1 SUMO (Suggested Upper Merged Ontology)

SUMO was originally an upper level ontology (Niles and Pease, 2001a), but now includes mid-level and domain ontologies as well (Pease, 2002). These domain-specific ontologies inherit the broad conceptual distinctions of SUMO, and specify the concepts and axiomatic content of a particular domain. Niles and Pease (2001b) explain SUMO's most general concepts and the relations between them, such as the distinction between `Object` and `Process`. Pease et al. (2002) illustrates the structure of SUMO and explains that the existence of SUMO-WordNet mapping serves as a gauge of the coverage of the ontology. Niles and Pease (2003) explain a SUMO-Wordnet mapping methodology and the specifics of the language used to show the synonymy, hypernymy, and instantiation relations.

In addition to the relations between entities that ontologies typically contain, SUMO provides relations between entities and processes, and statements in higher order logic (Benzmüller and Pease, 2012) that attempt to fully define each concept. For instance, the relation *result* is a binary relation with two arguments, a `Process` and an `Entity`, where the `Entity` is the output or product of the `Process`. So both temporally and causally, the `Process` is antecedent to the `Entity`. This relation is one of a few that help us find the agentive quale. Another relation between a `Process` and an `Entity` is *instrument*, where the `Entity` is used by an agent in bringing about a `Process`. So, in this case, the `Process` is consequent to the `Entity`. We use this relation to help find the telic quale for entities.

Relations are nodes under the `Abstract` hierarchy of SUMO, but many relations have sub-relations as well as instances. In general, since relations are like functions, each relation has a domain, and SUMO provides the set of arguments each relation can take. In some cases, the argument of a relation could be a whole formula by itself, so it does not point to a certain set. Rather, it is usually a complex axiom in itself, containing at least one Entity, at least one relation, and possibly some Processes. One of the main pointers we use to find telic quale is *hasPurpose*, which usually has such complex arguments. These arguments are themselves logical formulae, supported in the higher order logic employed for SUMO. For instance, SUMO stipulates the purpose of a Microwave using the *hasPurpose* relation, with a formula that can translate to: there is an entity PreparedFood and a process Heating, and the Microwave is an instrument for Heating, through which the PreparedFood gets heated.

3.2 REO (Rich Event Ontology)

The Rich Event Ontology (REO) unifies existing SRL schemas by providing an independent conceptual backbone through which they can be associated, and it augments the schemas with event-to-event causal and temporal relations. The ontology was developed in response to unsuccessful efforts to map directly between FrameNet (Fillmore et al., 2002) and the small set of disparate event types in Rich Entities, Relations and Events (ERE) (Song et al., 2015) (originally based on Automatic Content Extraction (ACE) (Doddington et al., 2004)). The difficulty was mainly due to differences in the granularity of events described by the FrameNet frames and ERE event types and inconsistencies in how the resources divided the semantic space.

FrameNet, ERE, and VerbNet (Schuler, 2005) have wide-coverage lexicons having to do with events, and they contribute annotated corpora and additional semantic and syntactic information that can be crucial to identifying events and their participants. REO serves as a shared hub for the disparate annotation schemas and therefore enables the combination of SRL training data into a larger, more diverse corpus, as well as expanding the set of lexical items associated with each ERE event type. By adding temporal and causal relational information, the ontology also facilitates reasoning on and across documents, revealing relationships between events that come together in temporal and causal chains to build more complex scenarios.

The structure of REO, illustrated in Figure 1, consists of a main "reference" ontology of generic event types and individual OWL "resource" ontologies. The relationships between the generic event types in the reference ontology and the event designations made in a particular annotated data resource are spelled out in various "linking models" that import both the reference ontology and a resource ontology. In the example shown, REO Discharge events map to the Releasing frame in FrameNet and the Release-Parole event type in ERE. However, other mappings between ERE and FN are necessarily more indirect. With respect to the Communication node in REO, ERE/ACE only maps to instrument_communication and statement while FN has mappings to nearly all daughters, as does VN, and VN maps to the mother Communication node as well. Additional indirect mappings are detailed in Brown et al. (2017). Individual words within the resource classes can be detected in text to find a wide variety of each event type, or one can query to view its participants and its relations to other events that is independent of the various lexical resource schemas.

4 Related Work

In considering how to best capture the type of information desired for relating ENDURANTS to PERDURANTS, we have examined the approaches of other benchmark ontologies. We find that the Basic Formal Ontology (BFO) (Smith et al., 2014) and the Descriptive Ontology for Linguistic and Cognitive Engineering (DOLCE) (Masolo et al., 2003) largely focus on parthood relations, qualities, and general notions of participation of an ENDURANT (or "continuant" as labeled in BFO) in a PERDURANT (or "occurrent" in BFO). The parthood relations have some overlapping information with the GL constitutive quale. Also of relevance to our work, BFO includes a class of what it calls "specifically dependent continuants," which includes functions, such as the function of a hammer to drive in nails.

The Event and Implied Situation Ontology (ESO) focuses on pre and post situations with respect to some event type, making explicit how a participant's pre and post states change in an event. Although there may be some amount of overlap in what ESO includes and the information captured in the GL agentive quale (e.g., the agentive quale for bread is baking), ESO's information will likely be much broader than the GL agentive, which essentially is limited to creation events. Although we hope to

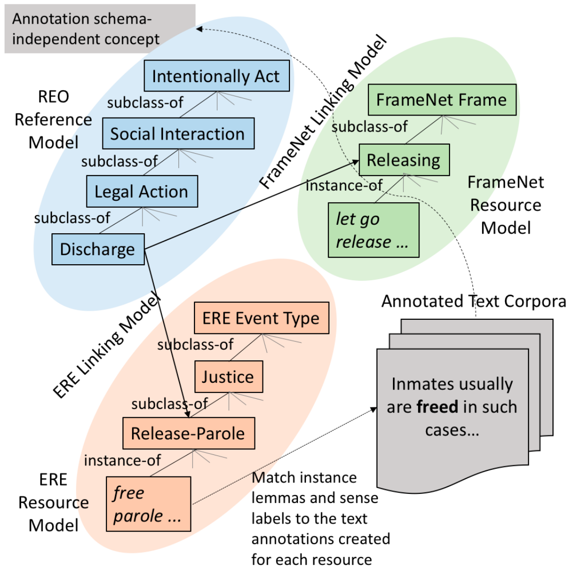

Figure 1: OWL resource models are linked to REO reference model in individual linking models (only ERE and FrameNet shown, but ACE and VerbNet linking models exist as well).

integrate or map the added and overlapping information from ESO, BFO and DOLCE in the future, here we focus on extracting information from SUMO.

5 Automatic Extraction of Qualia Relation Information

5.1 Methodology

In order to find qualia relations for entities in REO, we looked for ways to automatically extract them from SUMO. By examining around a hundred nodes in SUMO, a number of relations were found to be useful in extracting qualia. For instance, the relation *hasPurpose* in SUMO directly specifies the purpose (hence the telic quale) of a given entity. If a given entity is the second argument of the relation *instrument*, then it is a tool used to bring about an event which is the first argument of the relation *instrument* in the same axiom (e.g., (instrument ?TRANSFER ?ARTERY), where ARTERY is a tool enabling a TRANSFER event). Therefore, the first argument could be extracted as the telic quale of this entity. In the same way, a number of relations, including *part* and all its sub-relations, were found to lead us to the constitutive quale. For instance, we find that `Syllable` is a part of a `Word`, through an axiom roughly like (*part* `Syllable Word`). Though the sub-relations of *part* have fine-grained semantic differences, they all have the same weight here for our purposes. These include relations such as *initialPart, initiallyContainsPart, partTypes, typicalPart,* etc. The SUMO relation *result* leads us to the agentive quale, as its second argument is the output (or the product) of its first argument (e.g., (result ?WRITE ?TEXT) where TEXT is the output of WRITE event).

In addition to using SUMO relations, we used each node's documentation in SUMO to automatically extract telic quale using fairly straightforward regular expressions, since a number of entities were found in our initial examination to have their purpose described in their documentation, but not in their axioms. For instance, we extracted parts of the documentation that occur after terms such as 'purpose is,' 'intended to,' 'designed for,' etc. as potential candidates for the telic quale.

As the formal quale is basically an IS-A relationship, we extracted the parent of each node as its formal quale. Of course the inheritance is always at work for all the four types of qualia relations. For instance, an `Adverb` is a `Word` (parent), but is also a `LinguisticExpression` (grandparent).

The output of the function that automatically extracts qualia[1] currently can have two forms. If they are extracted from SUMO axioms, we have a set of functions to extract the actual SUMO node that

[1] `https://github.com/ghamzak/SUMOQualiaExtraction`

represents the desired quale. The second format of the output occurs when we extract qualia from the documentation. In this case, what we extract is actually plain English – a part of a sentence. In writing functions to extract sensible parts of sentences, we decided to sacrifice recall in favor of precision. At the end, we managed to extract 112 agentive, 762 telic, and 481 constitutive qualia relations. We also have a separate function that takes a SUMO entity as input and returns all possible qualia found in SUMO for that entity. For instance, `Building` has 13 total qualia found in SUMO, including 13 constitutive, 1 telic, and 1 agentive qualia relations. Table 1 illustrates the results of automatic extraction.

Agentive	112
Telic	762
Constitutive	481

Table 1: Number of Qualia found in SUMO using automatic extraction

5.2 Evaluation

In order to evaluate the automatically extracted qualia relations from SUMO, we designed an on-line annotation system, where we asked the annotators to decide whether the automatically extracted quale relation was reasonable or unreasonable, or indicate if they're not sure. We also had a comment box for each entry, and asked the annotators to provide comments. For instance, for the telic quale, our instruction for commenting was as follows:
If you feel the given function is reasonable but it's not at all how you would phrase or describe the function, then please provide your own description/phrasing of the function in the comment box.
You can also use to comment box to suggest a function for the unreasonable cases.
If you're unsure, comment on what makes you unsure.
In collecting such comments in addition to the reasonable judgments, we hope to gain insights into better alternatives for the quale relation.

We also provided SUMO documentation for the entity in question for cases where the annotator might not be familiar with the entity. The need for this was revealed in our pilot testing of the annotation system, where one of the annotators had not heard of some entities such as `Lanai`, which according to SUMO "refers to a roofed outdoor area Adjacent to a Building often furnished and used as a living room." So we decided to include SUMO documentation for each entity just in case the annotators are not familiar with it. Figure 2 shows a sample entry in our evaluation system.

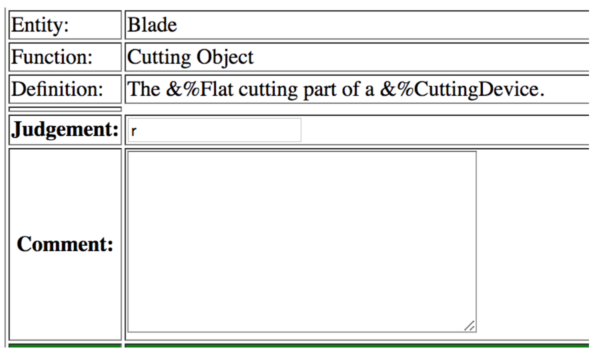

Figure 2: A sample entry in the evaluation system, asking about the telic quale.

To ensure that the annotators are not judging haphazardly and without thinking, we inserted 3 random attention tests in each page of annotation. The attention tests were completely unreasonable possibilities, in which the extracted function for one entity was paired with a totally different entity (e.g., Entity: AerobicExerciseDevice, Telic: to attach one thing to something else). So in each page, we had 25 real tasks and 3 attention tests. Examining the accuracy of each annotator, we had to throw out the data from one of them with only 73% accuracy on the attention tasks; those tasks were re-annotated by another annotator.

It would be prohibitively difficult to have a measure of recall for this task of automatic extraction, so we limit our evaluation to the precision of the results, using human annotators' judgments as the gold standard. Table 2 shows the precision measures found. The average interannotator agreement was 84.13%.

	Agentive	Telic	Constitutive
Reasonable	93%	88.26	85.89%
Unreasonable	2%	7.87%	10.74%
Unsure	5%	3.87%	3.37%

Table 2: Precision of Qualia Relations Extracted from SUMO

5.3 Error Analysis

The 'unreasonable' judgment for the constitutive relation was mostly applied to the ones taken from SUMO axioms, where the constitutive relation was too general for our purposes. For instance, the extraction found `Object` as a component part of `WireCoil`, and `Physical` as what constitutes `Solenoid`. Despite being true, they're too general to be accepted by a human as a reasonable part-whole relation. Another reason for marking the extracted constitutive relation as unreasonable was the jargon used in particular professions with which an annotators was not familiar, such as biology or chemistry. For instance, an `AtomicGroup` is part of a `Molecule`, but it's been marked as unreasonable with the comment: "part whole switch." Some other errors were due to an ordering mistake in SUMO axioms, such as (part Penne Hole) (which means Penne is part of Hole), whereas the reverse is true. Still others were due to a bug in the extraction, which ignored negation in axioms when finding constitutive relations, such as `BloodTypeB` which does NOT contain `AntigenA` according to SUMO axioms. We extracted it as a part by ignoring the negation. Thus, the results of the annotation were illuminating: helping to pinpoint where SUMO is too general for our purposes, or where our extraction script needs refinement.

The unreasonable judgments for the telic quale were mostly due to not yet capturing and combining inherited relations from SUMO. At any particular node, SUMO underspecifies the definition because it assumes inheritance from ancestor nodes. Although we assume this as well, the qualia relations we have extracted are limited to the ones found with direct mentioning of that entity. For instance, for the entity `MilitaryVehicle`, the telic role extracted was "MilitaryOrganization uses it," which was marked as unreasonable with the suggested alternative "Provide transportation for any military organization." However, `MilitaryVehicle` has `Vehicle` as its parent and inherits from it. Therefore, the telic quale for `Vehicle`, which is extracted as "Translocation," would be inherited by `MilitaryVehicle`.

This sort of error confirms our plan to inherit qualia relations and add the quale for each entity to all its children. Not only will these types of errors be eliminated, the number of entities with qualia relations will increase significantly. Currently, many lower level entities have no specific qualia to be extracted at their SUMO node but have very informative quale that could be inherited from parent nodes. Thus, we would have a significant increase in coverage (recall), while precision is guaranteed to remain high.

Yet other instances of unreasonable telic quale were due to the wording used in SUMO. For instance, the following pairs have been judged as unreasonable.
Entity: HearingProtection – Function: protect Human from Injuring caused by RadiatingSound
Entity: PerformanceStage – Function: location of Demonstrating
Entity: Campground – Function: to have MobileResidences
These may not be how a human would describe the functions, but SUMO has tried to maximize the grounding of its definitions by using its other defined entities within them, leading to a more interconnected network of concepts. `Demonstrating`, for instance, is not the word people use to talk about the function of a performance stage, but according to the documentation of `Demonstrating` in SUMO, it would cover 'software demos, theatrical plays, lectures, dance and music recitals, museum exhibitions,

etc.' Given the connection of REO concepts to SUMO, we need not be overly concerned about these types of seemingly inaccurate results.

6 Incorporating Qualia Relations into REO

6.1 Expanding REO Endurants

As an ontology meant to provide a shared vocabulary of SRL annotation resources, the development of REO has focused on how to structure ontological relations between the events and states included in these resources (ACE, ERE, VerbNet, FrameNet). Although our intention has never been to create a detailed ontology of the entities that serve as the participants in these events/states, such information is needed to generalize and map this information across resources, as well as provide some insights into selectional restrictions generally and qualia relations specifically. Thus, we opted to integrate REO with an existing ontology containing the type of information on participants that we were interested in.

Given the detailed information on objects and accompanying information such as *hasPurpose* in SUMO, we decided that it was best suited to our needs. REO aligns with early distinctions made by DOLCE, distinguishing between ENDURANT Entities and PERDURANT Entities, where ENDURANTS are defined as "Those entities that can be observed/perceived as a complete concept, no matter which given snapshot of time," whereas PERDURANTS are defined as "Those entities for which only a part exists if we look at them at any given snapshot in time. Variously called events, processes, phenomena, or activities and states, PERDURANTS have temporal parts or spatial parts and participants." PERDURANTS largely map to what SUMO deems `Processes`, and although many ENDURANTS map to what SUMO deems `Objects`, there were also a variety of useful concepts in SUMO that seemed to fit the definition of an ENDURANT. These include: `ContentBearingPhysical`, `FinancialAsset`, `Graph`, `GraphElement`, `Model`, `PhysicalSystem`, `Quantity`, and `SetOrClass`. To flesh out the ENDURANT portion of REO, we opted to integrate all daughters of `Object` and the classes just mentioned as daughters of our ENDURANT class. Specifically, this was done in a linking model, similar to the linking models that serve to map the generic event concepts in REO's 'Reference Model' to particular groups of events in the SRL resource models (see Brown et al. (2017) for more details on the structure of REO). The linking model imports an abbreviated OWL version of SUMO (containing only the ENDURANT-compatible classes mentioned, their documentation that provides a description of what they are, and their associated WordNet sense keys), and the REO OWL reference model. The resulting model has all of the REO event ontology nested under PERDURANT, and the extracted SUMO content nested under ENDURANT (see Figure 3).

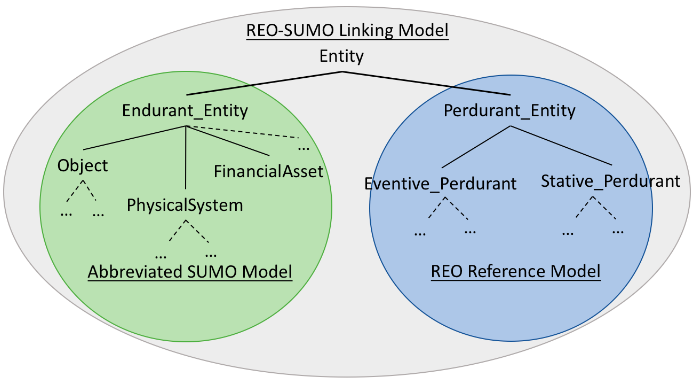

Figure 3: REO-SUMO Linking Model: This model imports the abbreviated SUMO model (containing ENDURANT-type classes only) and the REO reference model. Thus, REO specifies the PERDURANTS of the model, while SUMO specifies the ENDURANTS.

Admittedly, there are a variety of weaknesses to this model. A model that used all of SUMO would contain the many axioms found in the full version of SUMO, which OWL is not expressive enough to handle (mathematically a description logic cannot state the content of first order logic, much less in this case, higher order logic). These axioms facilitate some of the rich reasoning potential of SUMO, which is lost in our more superficial model. We are therefore considering new ways in which to integrate the lexical resource information included in REO into the full SUMO model, enabling users to take advantage of SUMO with respect to reasoning about events included in the SRL annotations in the future.

6.2 Initial Integration of Qualia Relations

As a preliminary step in incorporating the extracted qualia relations into our ontology, we have simply added the information as annotations on the ENDURANTS drawn from the abbreviated SUMO model. We have introduced the following annotation properties: has_Quale_Agentive, has_Quale_Constitutive, and has_Quale_Telic. We do not include an explicit relation for the Formal quale, as we assume this information is encoded in the ontological structure itself. A variety of examples of these qualia incorporated are given in Table 1.

Endurant	**has_Quale...**	**Concept**
Brake	Telic	Decelerating
Building	Constitutive	Wall
Organization	Agentive	Founding

Table 3: Examples of has_Quale annotations of Endurant entities.

More general parent classes lack cohesive qualia. In fact, almost none of the nine direct daughters of ENDURANT (e.g., `Object`, `PhysicalSystem`, `Model`, etc.) have qualia annotations, but the daughters of these classes (and subclasses therein) are increasingly populated with qualia annotations, reflecting their increasing specificity.

In future work, we will incorporate the extracted qualia relation information (along with the comments acquired through annotators applied to them) into the REO/SUMO Linking Model using object properties (as opposed to the current annotation properties). These object properties will relate SUMO ENDURANTS to either REO PERDURANTS or other SUMO ENDURANTS, depending upon the relation, such that the ontology can be exploited for telic, agentive, formal, and constitutive qualia, as well as the events in which an endurant is typically involved. For example: `Meal` (ENDURANT extracted from SUMO) has_Quale_Telic CONSUMPTION (PERDURANT from REO). Unlike the annotation properties currently implemented, the object properties will better facilitate querying and reasoning using the ontology, and will effectively map the concepts extracted from SUMO (e.g., an extracted telic quale for `meal` is 'contains nutrients for humans') to the event classes in REO (labeled CONSUMPTION). We plan to do this mapping manually using a group of expert annotators who are familiar with the ontology and the linked lexical resources.

In addition, we plan to augment ongoing research evaluating the utility of REO in scene understanding in images and video (Tahmoush and Bonial, 2016). Specifically, we will be exploring how the information relating ENDURANTS and PERDURANTS for a given event type might be leveraged in putting together the pieces relating objects recognized in an image to higher-order scene understanding and perhaps even activity recognition. This is inspired in part by work on "visual semantic role labeling" using FrameNet for situation understanding (Yatskar et al., 2016), but we expect to exploit the more general qualia information to understand, for example, what type of room is encountered if a bed is recognized, versus a tooth brush, based on what we know about how those objects are typically used.

7 Conclusion

An understanding of events is not complete without understanding their participants. Similarly, an understanding of objects is incomplete without some knowledge of what events they are typically associated with. We have described efforts to extract GL qualia relations using a novel methodology exploiting particular information from SUMO relations and documentation. Our evaluation shows that the vast majority of the automatically extracted relations are judged reasonable by human annotators. Our motivation has been to enrich REO with origin, function and part-whole information in the form of GL qualia relations. To lay the groundwork for this effort, we have first integrated REO with ENDURANTS drawn from carefully selected SUMO classes. We have completed an initial integration of the extracted qualia relation information as annotations in the updated REO. This version of the ontology will be made freely available, as will subsequent versions of the ontology that incorporate the qualia relations as object properties.

References

Christoph Benzmüller and Adam Pease. 2012. Higher-order aspects and context in sumo. *Web Semantics: Science, Services and Agents on the World Wide Web*, 12:104–117.

Susan Brown, Claire Bonial, Leo Obrst, and Martha Palmer. 2017. The rich event ontology. In *Proceedings of the Events and Stories in the News Workshop*, pages 87–97.

George R Doddington, Alexis Mitchell, Mark A Przybocki, Lance A Ramshaw, Stephanie Strassel, and Ralph M Weischedel. 2004. The automatic content extraction (ace) program-tasks, data, and evaluation. In *LREC*, volume 2, pages 837–840.

Charles J Fillmore, Collin F Baker, and Hiroaki Sato. 2002. The framenet database and software tools. In *LREC*.

Alessandro Lenci, Nuria Bel, Federica Busa, Nicoletta Calzolari, Elisabetta Gola, Monica Monachini, Antoine Ogonowski, Ivonne Peters, Wim Peters, Nilda Ruimy, et al. 2000. Simple: A general framework for the development of multilingual lexicons. *International Journal of Lexicography*, 13(4):249–263.

Beth Levin. 1993. *English verb classes and alternations: A preliminary investigation*. University of Chicago press.

Claudio Masolo, Stefano Borgo, Aldo Gangemi, Nicola Guarino, and Alessandro Oltramari. 2003. Wonderweb deliverable d18, ontology library (final). *ICT project*, 33052.

Ian Niles and Adam Pease. 2001a. Origins of the ieee standard upper ontology. In *Working notes of the IJCAI-2001 workshop on the IEEE standard upper ontology*, pages 37–42.

Ian Niles and Adam Pease. 2001b. Towards a standard upper ontology. In *Proceedings of the international conference on Formal Ontology in Information Systems-Volume 2001*, pages 2–9. ACM.

Ian Niles and Adam Pease. 2003. Mapping wordnet to the sumo ontology. In *Proceedings of the ieee international knowledge engineering conference*, pages 23–26.

Adam Pease, Ian Niles, and John Li. 2002. The suggested upper merged ontology: A large ontology for the semantic web and its applications. In *Working notes of the AAAI-2002 workshop on ontologies and the semantic web*, volume 28, pages 7–10.

Adam Pease. 2002. Sumo.

Adam Pease. 2011. *Ontology: A practical guide*. Articulate Software Press.

James Pustejovsky. 1991. The generative lexicon. *Computational linguistics*, 17(4):409–441.

Karin Kipper Schuler. 2005. Verbnet: A broad-coverage, comprehensive verb lexicon.

Barry Smith, Mauricio Almeida, Jonathan Bona, Mathias Brochhausen, Werner Ceusters, Melanie Courtot, Randall Dipert, Albert Goldfain, Pierre Grenon, Janna Hastings, et al. 2014. Basic formal ontology 2.0 draft specification and user's guide.

Zhiyi Song, Ann Bies, Stephanie Strassel, Tom Riese, Justin Mott, Joe Ellis, Jonathan Wright, Seth Kulick, Neville Ryant, and Xiaoyi Ma. 2015. From light to rich ere: annotation of entities, relations, and events. In *Proceedings of the 3rd Workshop on EVENTS at the NAACL-HLT*, pages 89–98.

David Tahmoush and Claire Bonial. 2016. Ontology-based improvement to human activity recognition. In *Automatic Target Recognition XXVI*, volume 9844, page 98440U. International Society for Optics and Photonics.

Mark Yatskar, Luke Zettlemoyer, and Ali Farhadi. 2016. Situation recognition: Visual semantic role labeling for image understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5534–5542.