Exploiting Metonymy from Available Knowledge Resources

Itziar Gonzalez-Dios(✉), Javier Álvez, and German Rigau

IXA and LoRea Groups, University of the Basque Country UPV/EHU, Donostia/San Sebastian, Spain
{itziar.gonzalezd,javier.alvez,german.rigau}@ehu.eus

Abstract. Metomymy is challenging for advanced natural language processing applications because one word or expression is used to refer to another implicit related concept. So far, metonymy has been mainly explored in the context of its recognition and resolution in texts. In this work we focus on exploiting metonymy from existing lexical knowledge resources. In particular, we analyse how metomynic relations are implicitly encoded in WordNet and SUMO. By using an existing automated reasoning framework to test the new modelling acquired from WordNet and SUMO, we propose a practical way to deal with figurative expressions.

Keywords: Metonymy · WordNet · SUMO · Knowledge representation · Commonsense reasoning

1 Introduction

Metonymy is a cognitive and linguistic process where one thing is used to refer another [1]. Together with the metaphor, metonymy is an imaginative device (a figure of speech) that is an object of prime interest for cognitive linguistics. In the conceptual projection of metonymy, the target, which belongs to an experiential domain, is understood in terms of the source, which belongs to other domain [2]. Common metonymic cases are expressing parts as wholes, continents for contents, places for products... For example, in (1) the word *Bordeaux* is not referring to the French city, but to the wine produced in its region. That is, a place is used for a product.

(1) I never drink a young Bordeaux.

The use of figurative language is crucial for natural language understanding and, therefore, for Natural Language Processing (NLP). In the case of metonymy, many data-driven works have dealt with its resolution. For example, Markert et al. [3] present the results of the SemEval-2007 shared task on that topic and Gritta et al. [4] focus on locations. There are also some works such as [5,6] which investigate logical metonymy, an elliptical construction that occurs when an event-subcategorising verb is combined with an entity-denoting direct object.

A. Gelbukh (Ed.): CICLing 2019, LNCS 13451, pp. 34–43, 2023.
https://doi.org/10.1007/978-3-031-24337-0_3

In this work we analyse the characterisation of metomymy in existing Lexical Knowledge Resources (LKRs). The LKRs we use in this paper are the semantic network WordNet [7] and the SUMO ontology [8]. WordNet [7] has implicit metonymic relations that need to be extracted and SUMO has no formal representation of metonymy. That is why, we explore how we can reuse the information in WordNet to model it in SUMO, and start a discussion towards making these relations explicit and systematised in WordNet. We concentrate on the membership relation of locations and organisations that are coded in WordNet. For instance, the words *Andorra* and *Hanseatic League* denote respectively in (2) and (3) a government that is formed by group of people and a group of cities formed by group of people, and not as regions or organisations on their own, but we use them with metonymic reading to get a more efficient and economic communication.

(2) <u>Andorra</u> does not close border with Catalonia.

(3) <u>The Hanseatic League</u> carried out an active campaign against pirates.

If we look for Andorra and Hanseatic League in WordNet we do find that Andorra (nation) has Andorran (inhabitant) as member and Hanseatic League (political organisation) has Bremen (city) as its member. That is, Andorra and Hanseatic League are represented also as group of people in WordNet and not only as locations or organisations. That is why, we think that the member relation in WordNet can be a good clue to detect metonymic readings.

So, the aim of this paper is to detect metonymic candidates via the member relation in WordNet and propose an approach to show that it is possible to deal with figurative expressions. Exactly, we explore how we can reuse the information in WordNet to model it in SUMO, and, moreover, we propose to start a discussion towards making these relations explicit and systematised in WordNet. Another aim of this work is to ease the burden and the cost of creating new resources by exploring the interoperability of the ones that are available in order to create a future framework for commonsense reasoning.

This paper is structured as follows: in Sect. 2 we present the analysis of metonymy in the LKRs, in Sect. 3 we explain our approach to treat metonymy that we evaluate in Sect. 4, we discuss the results in Sect. 5 and we conclude and outline the future work in Section 6.

2 Analysing Metonymy in WordNet and SUMO

Metonymic readings in WordNet are not specifically labelled, but some are coded within the *member* relation. Based on the annotation framework of [3] where they distinguished metonymic patterns for locations (place-for-people, place-for-event...), organisations (org-for-members, org-for-event...) and class-independent categories (object-for-name and object-for-representation), we have decided to extract the synsets denoting locations and organisations, which, in our opinion, are candidates for being metonymic.

Table 1. Selected BLCs to find metonymic candidates

Candidate type	Selected BLCs
Locations	$area^1_n$, $capital^3_n$, $district^1_n$, $country^2_n$, $state^1_n$, $American\_State^1_n$, $geographical\_area^1_n$, $desert^1_n$, $city^1_n$, $national\_capital^1_n$, $state\_capital^1_n$, $national\_park^1_n$, $town^1_n$, $region^1_n$, $boundary^1_n$, end^1_n, $port^1_n$
Organisations	$organization^1_n$, $gathering^1_n$, $lineage^1_n$, $terrorist\_organization^1_n$, $artistic\_movement^1_n$

In order to extract the candidates, we have selected the Basic Level Concepts (BLCs) [9] presented in Table 1 relating locations and organisations. In this paper, we will use the format $word^s_n$, where s is the sense number of the *word* and n means that it is a noun. For brevity, we will only show a representative variant of each selected BLC synset in the table. Based on those BLCs we have also gathered their hyponyms as long as their semantic file was *location* for locations and *group* for organisations. In total we have selected 156 candidate synsets.

Following, among those candidates we have extracted the synset pairs that are linked via the *member* relation. This way, we have obtained a dataset of 672 metonymic pairs. This way, we have detected some of the implicit metonymic readings in WordNet.

The SUMO counterpart of the WordNet *member* relation is the predicate $member_r$ which relates an individual object (i.e. an instance of the SUMO class $Object_c$, the class for ordinary objects) as part of a collection (i.e. an instance of $Collection_c$, the class where its member can be altered without changing the nature of the collection itself). To denote SUMO concepts in this paper we will use the symbol c for SUMO classes and r for SUMO relations.

Moreover, WordNet and SUMO are linked via a semantic mapping [10]. This mapping links a synset to a SUMO concept by means of the relations *equivalence* (both mean the same) or *subsumption*/*instantiation* (the semantics of the SUMO concept is more general than the semantics of the synset in WordNet). The semantic mapping relation is denoted, in this paper, by appending as suffix the symbols '=' (equivalence), '+' (subsumption) and '@' (instantiation) to the corresponding SUMO concept. This way, the synset $Hardy^1_n$ presented in (4) is mapped to Man_c+ and $Laurel\_and\_Hardy^1_n$ is to $Group_c+$. It is important to mention that no metonymic readings are coded in the original mapping, only literal.

So, if we cross-check the knowledge in WordNet (the pairs in the dataset we have extracted) and SUMO (see Sect. 4), we see that example (4) is validated: from the knowledge in SUMO it is possible to infer that some instances of Man_c and $Group_c$ are related be $member_r$. In other words, humans can be members of groups. On the contrary, Example (5) is unvalidated since $France^1_n$ is connected to $Nation_c+$ and $Nation_c$ is not instance of $Collection_c$. That is, $France^1_n$ is not understood as a group of people in the mapping to SUMO, but as a mere region (an object) and, therefore, cannot have members.

(4) a. $Hardy^1_n$ United States slapstick comedian who played (...) (Man_c+)
—*MEMBER OF*—
b. $Laurel\_and\_Hardy^1_n$ United States slapstick comedy duo (...) ($Group_c$+)

(5) a. $French\_person^1_n$ a person of French nationality. ($EthnicGroup_c$+)
—*MEMBER OF*—
b. $France^1_n$ a republic in western Europe; (...) ($Nation_c$+)

Once we have extracted the metonymic candidates, we pursue to model metomymic readings, but where should it be modelled? In the mapping? In the ontology? In both? In the Sect. 3 we present our practical approach to deal with this figure of speech.

3 Modelling Metonymy

In order to model metonymy based on the WordNet dataset we have analysed both the mapping and the ontology.

3.1 Exploring the Mapping

By inspecting the mapping manually by class frequency, we have realised that there is lack of systematicity above all with demonyns: for example (6) and (7) are mapped to $FreshWaterArea_c$+ (the class for rivers and lakes) and $NaturalLanguage_c$+ (the class for languages) respectively although each one is an individual, (8) is mapped to $EthnicGroup_c$+ (the class for groups of people that share a country, a language...) despite being also a person. The latest is the most frequent case in the dataset.

(6) $Sabine^2_n$ a member of an ancient Oscan-speaking people (...)

(7) $Sotho^1_n$ a member of the Bantu people who (...)

(8) $Sherpa^1_n$ a member of the Himalayan people (...)

So, in order to correct automatically the mapping of synsets like (8), we have created the following heuristic based on the semantic files (SF), the glosses of WordNet the BLCs:

- Change the mapping of a synset from $EthnicGroup_c$ to $Human_c$ (the class for modern man and woman) if its BLC is $person^1_n$; its SF is *person*; and there is one of the following expressions *a member of, a native or inhabitant of, an inhabitant of, a person of, a speaker of, a German inhabitant of, a native or resident or, a resident of, a native of, a Greek inhabitant of, a native or resident of, a person of, inhabitant of, a Polynesian native or inhabitant of* or *an American who lives in* in its gloss.

This way, we have corrected the mapping of 272 synsets.

We have also manually analysed the mapping of the synsets that are connected to a SUMO concept different from $EthnicGroup_c$ but that fulfil the remaining above conditions. Out of 75 synsets we have corrected 30, most of them corrected to the SUMO class $Human_c$, namely cases such as (6) and (7).

For the other kind of errors (not relating demonyns), we have extracted the synsets containing the expressions *an island territory of, an island in* and *a glacial lake*. Out of 6 synsets we have corrected 3 manually.

In total, we have corrected the mapping 308 synsets: 33 manually and 272 automatically.

After having corrected the mapping, we have modelled the metonymy. To that end, we have added the $Group_c$+ class to the synsets in the *b* part of the pair. This way, we have created a corrected and modelled mapping (CMM).

3.2 Exploring the Ontology

In order to start exploring the ontology we have performed an error analysis on the most frequent metonymic pairs that remained unresolved in [11]. The first inspection revealed us the respective members of cities, countries... and groups or organisations as members of other groups or organisations were not validated. As mentioned before, the predicate $member_r$ in SUMO relates instances of $Object_c$ and $Collection_c$, but regions, for example, are not connected to (subconcepts of) $Collection_c$ and groups are not connected to (subconcepts of) $Object_c$. So, according to SUMO, they cannot be related as part or whole in member pairs.

After this analysis, we have included 22 new general axioms in the ontology for organisations and regions that modelled metonymy and added missing related knowledge. Next, we describe two of the introduced axioms:

– Relating the regions, any SUMO concept related to some instance of $LandArea_c$ (the general class for areas where predominates solid ground and includes all the nations, states, provinces...) by $member_r$ as whole is restricted to be instance of $Agent_c$ (the class for something or someone that can act on its own and produce changes), which is superclass of both $Organization_c$ (the class for corporations or similar institutions) and $Human_c$. That means that countries can have humans and organisations as members. This way, we will validate example (2): $Andorra_n^1$ ($Nation_c$+) can have $Andorran_n^1$ ($Human_c$+) as member.

```
(forall (?LAND ?MEMBER)
    (=>
        (and
            (instance ?LAND LandArea)
            (member ?MEMBER ?LAND))
        (instance ?MEMBER ?Agent)))
```

– Relating the organisations, any instance of $PoliticalOrganization_c$ (the class for organisation that attempt some sort of political change) is restricted to be related by $member_r$ as whole to some instance of $GeographicArea_c$ (the general class for geographic locations with definite boundaries). In other words, political organisations can have countries, cities... as members. This way, we will validate example (3): $Hanseatic\_League^1_n$ ($PoliticalOrganization_c$+) can have $Bremen^1_n$ ($City_c$+) as member.

```
(forall (?PORG)
    (=>
        (instance ?PORG PoliticalOrganization)
        (exists (?GAREA)
            (and
                (instance ?GAREA GeographicArea)
                (member ?GAREA ?PORG)))))
```

With this kind of axioms we look also for the validation of candidates such as $Breton^1_n$ as member of $Bretagne^1_n$ ($StateOrProvince_c$+) and $Belgique^1_n$ ($Nation_c$+) as member of $Benelux^1_n$ ($PoliticalOrganization_c$+).

4 Evaluation Framework and Results

In order to evaluate our interventions, we have used the framework presented by [11,12], where competency questions (CQ) based on predefined question patterns (QP) are automatically created taking the knowledge of the WordNet dataset into account. In Fig. 1 we present an example of the CQs that are created.

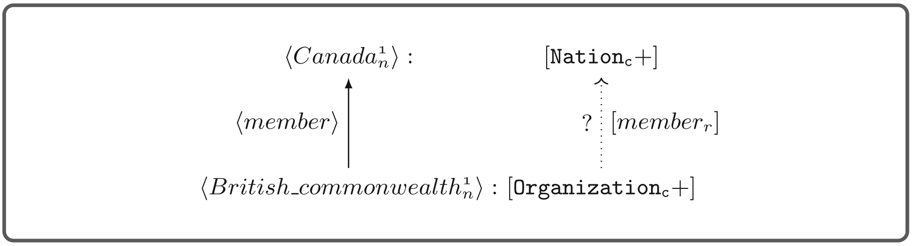

Fig. 1. Example of competency question about metonymy relating countries and organizations

In the present work, we have used the QPs presented in the above mentioned work in order to create the CQs for the *member* relation with the original mapping (OM), the original ontology (OO), our corrected and modelled mapping

(CMM) and our modelled ontology (OM). In order to see which intervention has more impact, we have considered the following cases:

1. OM+OO: The starting point with the original mapping and original ontology
2. CMM+OO: Measuring the impact of the mapping with the corrected and modelled mapping and original ontology
3. OOM+MO: Measuring the impact of the ontology with the original mapping and modelled ontology
4. CMM+MO: Measuring the impact of both the mapping and the ontology with the corrected and modelled mapping and modelled ontology

In the first and third cases, 89 CQs are created, while 124 CQs are created in the second and fourth cases. Then, the resulting sets of CQs are experimentally evaluated by using the ATP Vampire v4.2.2 [13] in a Intel® Xeon® CPU E5-2640v3@2.60GHz with 2GB of RAM memory per processor. For each test, we have set an execution-time limit of 300 s and a memory limit of 2GB.[1] Depending on the outcome provided by the ATP, each CQ can be classified as i) *passing* (the ATP proves that the knowledge encoded in the CQ is entailed by SUMO); ii) *non-passing* (the knowlegde encoded in the CQ is incompatible with SUMO since the ATP proves that its negation is entailed); or iii) *unknown* (if the ATP finds no proof).

The results of these experiments are presented in Table 2. More concretely, we provide the number of metonymic pairs that are validated (Validated column), unvalidated (Unvalidated column) and unclassified (Unclassified column). Roughly speaking, a metonymic pair is classified as validated/unvalidated/unclassified if the corresponding CQ is decided to be passing/non-passing/unknown respectively. Further, a metonymic pair is also classified as unvalidated if the mapping information is incompatible with the SUMO predicate $member_r$, since no CQ can be obtained. We provide the number of metonymic pair with incompatible mapping information in the column Unvalidated between brackets. On this basis, we calculate the recall (Recall column), precision (Precision column) and the $F1$ measure (F1 column): recall is validated pairs divided by total pairs, precision is validated pairs divided by the sum of validated and unvalidated pairs and $F1$ is the harmonic mean between precision and recall.

Table 2. Experimental results

Experiment	Validated	Unvalidated		Unclassified	Recall	Precision	$F1$
OM+OO	4	357	(356)	311	0.006	0.011	0.008
CMM+OO	5	12	(11)	655	0.007	0.294	0.015
OM+MO	197	359	(356)	116	0.293	0.354	0.321
CMM+MO	483	12	(11)	177	0.719	0.976	0.828

[1] Parameters: `--proof tptp --output_axiom_names on --mode casc -t 300 -m 2048`.

Looking at the experimental results (Table 2), we see that in the first case (OM+OO) only 4 pairs are validated. Thus, the $F1$ is really low (0.008). By correcting and modelling the mapping (CMM+OO), only one more pair is validated, and although the $F1$ increases the (0.015), the results are still very poor. The important comment on this experimental phase is that more pairs are unclassified (345 more) and that they are no longer unvalidated. That is, this time the mapping information is not incompatible with $member_r$, but the knowledge encoded in the corresponding CQ is not yet entailed by the ontology. When adding metonymic knowledge in the ontology, but still using the original mapping (OM+MO), we see that there is a big improvement in the number of validated pairs: 193 more pairs are validated (total 197) and, thus, $F1$ is increased if compared to the two previous cases. Further, if we compare the results in the cases CMM+OO and OM+MO, we see that modelling the metonymic knowledge in the ontology has more impact than modelling it in the mapping. Finally, by combining both mapping and ontology modellings (CMM+MO), we see that the results are undoubtedly better: 483 pairs are validated (more than the two thirds of the pairs) and $F1$ is 0.828, i.e. 0.820 more than in the initial case. However, there is room for improvement: 177 pairs are unclassified and in the case of 12 pairs the mapping information is not compatible with $member_r$.

In order to look at the results qualitatively, we present two validated examples: (9) and (10). Like (2), (9) is validated now because we added axioms that allow countries, regions... to have people and regions as members. Similar to (3), (10) is validated now because we added axioms to relate countries, nations... to organisations.

(9) a. $Mexican^1_n$ a native or inhabitant of Mexico (...) ($Human_c$+)
—*MEMBER OF*—
b. $Mexico^1_n$ a republic in southern North America(...) ($Nation_c$+)

(10) a. $France^1_n$ a republic in western Europe (...) ($Nation_c$+)
—*MEMBER OF*—
b. $NATO^1_n$ an international organization (...) ($Organization_c$+)

5 Discussion

As mentioned in the previous sections, we have a big improvement, but there still room to work with. That is why we have also performed a detailed analysis of the unvalidated and unclassified pairs.

Regarding the unvalidated pairs, only one pair with compatible mapping information is unvalidated: $crew^2_n$ (connected to $GroupOfPeople_c$+) as member of $workforce^1_n$ (connected to $SocialRole_A$+). The reason is that the class of attributes $SocialRole_A$ can be applied to only individuals, thus not to groups. Being the only one, this case was part of the long tail that we did not analyse during the inspection of the mapping. The other unvalidated pairs have incompatible mapping information: one of the synsets is mapped to a process and therefore

CQs cannot be created e.g. *conferee*$_n^2$ (connected to *Human*$_c$+) as member of *conferee*$_n^1$ (connected to *FormalMeeting*$_c$+, which is a SUMO process).

Regarding the unclassified pairs, we need to keep on modelling specific data, as we only modelled general metonymic relations in the ontology. We have also detected mapping errors in 4 synsets whose expressions did not cover such as *an Oscan-speaking member of*. Moreover, we have identified at least 38 pairs could be validated if we gave the ATP more time.

On the other hand, while modelling the information in the ontology we have realised that the *member* relation is not always clear in WordNet. Is an organisation *member of* another organisation? Or is it *part of* it? Does it depend? We think that sometimes the distinction is not very clear. For example, the synset *workforce*$_n^1$ has a member which is *crew*$_n^2$ (an organized group of workmen) and a part which is *shift*$_n^7$ (a crew of workers who work for a specific period of time).

That is why we consider that a discussion on the WordNet relations can be necessary. In our opinion, there is a lack of formality in the relations: some of them are ambiguous. So, we propose to start thinking about systematising them, axiomatising them. Moreover, we think that making explicit i.d. modelling implicit or hidden relations such as the case of metonymy we have presented here can be advantageous and profitable for NLP applications.

6 Conclusion and Future Work

In this paper we have presented a practical approach to show that it is feasible to work with figurative language in well-known knowledge resources by reusing their information and correcting their discrepancies. We have detected implicit metonymic relations regarding organisations and locations in WordNet and we have modelled this knowledge in SUMO. We have explored two general formalisations: modelling in the mapping and modelling in the ontology, and we have seen the changes in the ontology have more impact than the changes in the mapping. Moreover, we have corrected discrepancies in the mapping between them. As for the results, we see that there is considerable improvement (from $F1$ 0.008 to $F1$ 0.828) when merging both formalisations, but there is still room for improvement by given the ATP more resources or continuing adding more specific information. But, we have also seen that relations are not well-defined.

Therefore, future work involves also examining more precisely the meaning and the coherence of the relations both in WordNet and SUMO in order to systematise and axiomatise them. Moreover, we would like to study other kind of metonymic relations, for instance, relating processes and events. Evaluating the correctness of the implicit metonymic knowledge in WordNet is also one of our future goals.

Acknowledgments. This work has been partially funded by the Spanish Projects TUNER (TIN2015-65308-C5-1-R) and GRAMM (TIN2017-86727-C2-2-R) and the Basque Projects BERBAOLA (KK-2017/00043) and LoRea (GIU18/182).

References

1. Littlemore, J.: Metonymy. Cambridge University Press, Cambridge (2015)
2. Barcelona, A.: Metaphor and Metonymy at the Crossroads: A Cognitive Perspective. Walter de Gruyter, Berlin (2012)
3. Markert, K., Nissim, M.: Data and models for metonymy resolution. Lang. Resour. Eval. **43**(2), 123–138 (2009). https://doi.org/10.1007/s10579-009-9087-y
4. Gritta, M., Pilehvar, M.T., Limsopatham, N., Collier, N.: Vancouver welcomes you! minimalist location metonymy resolution. In: Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 1248–1259 (2017)
5. Shutova, E., Kaplan, J., Teufel, S., Korhonen, A.: A computational model of logical metonymy. ACM Trans. Speech Lang. Process. (TSLP) **10**(3), 1–28 (2013)
6. Chersoni, E., Lenci, A., Blache, P.: Logical metonymy in a distributional model of sentence comprehension. In: 6th Joint Conference on Lexical and Computational Semantics (* SEM 2017), pp. 168–177 (2017)
7. Fellbaum, C. (ed.): WordNet: An Electronic Lexical Database. MIT Press, Cambridge (1998)
8. Niles, I., Pease, A.: Towards a standard upper ontology. In: Guarino, N., et al., (ed.): Proceedings of the 2nd International Conference on Formal Ontology in Information Systems (FOIS'01). ACM, pp. 2–9 (2001)
9. Izquierdo, R., Suárez, A., Rigau, G.: Exploring the automatic selection of basic level concepts. In: Proceedings of the International Conference on Recent Advances on Natural Language Processing (RANLP'07), vol. 7 (2007)
10. Niles, I., Pease, A.: Linking lexicons and ontologies: mapping wordnet to the suggested upper merged ontology. In: Arabnia, H.R., (ed.): Proceedings of the IEEE International Conference on Information and Knowledge Engineering (IKE'03). vol. 2, pp. 412–416. CSREA Press (2003)
11. Álvez, J., Gonzalez-Dios, I., Rigau, G.: Cross-checking using meronymy. In: Calzolari, N., et al., (eds.): Proceedings of the 11th International Conference on Language Resources and Evaluation (LREC'18). European Language Resources Association (ELRA) (2018)
12. Álvez, J., Lucio, P., Rigau, G.: A framework for the evaluation of -based ontologies using. IEEE Access **7**, 1–19 (2019)
13. Kovács, L., Voronkov, A.: First-order theorem proving and Vampire. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 1–35. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-39799-8_1