Using Referential Language Games for Task-oriented Ontology Alignment *

** Nikolaos Kondylidis[0000-0003-4304-564X]
nikos.kondylidis@vu.nl

Vrije Universiteit Amsterdam, Amsterdam, The Netherlands

Abstract. Ontology Alignment (OA) is generally performed by requesting two parties to provide their complete knowledge to a third party that suggests potential schema alignments. This might however not always be possible or helpful, as for example, when two organisations want to query each other's knowledge, and none of them is willing to share their schema due to information privacy considerations. This Ph.D. explores how to allow multi-agent communication in cases where agents operate using different ontologies that cannot be fully exposed or shared. Our preliminary experiments focus on the case where agents' knowledge is describing a common set of entities and has the form of Knowledge Graphs (KGs). The suggested methodology is based on the grounded naming game, where agents are forced to develop their own language in order to refer to corresponding schema concepts of different ontologies. This way, agents that use different ontologies can still communicate successfully for a task at hand, without revealing any private information. We have performed some proof of concept experiments applying our suggested method on artificial cases and we are on the process of extending our methodology so that it can be applied in real-world KGs.

Keywords: Multi-Agent Communication · Task-oriented Ontology Alignment · Instance-based Ontology Matching.

1 Introduction and Motivation

A populated ontology is attributing characteristics to a set of instances using its ontology schema. Different schema designs and characteristics can be used for describing the same instances, depending on the purpose of the ontology. Even if these ontologies have common characteristics, these might not be communicated directly as it is expected to be defined under different symbols or namespaces. Ontology Alignment (OA) techniques attempt to bridge this gap by providing symbol alignments across ontologies, denoting that the meaning of these symbols is equivalent. Provided these alignments, the ontologies are able to represent their knowledge in terms of both ontologies, allowing them to query one another.

* This work was supported by "MUHAI - Meaning and Understanding in Human-centric Artificial Intelligence" project, funded by the European Union's Horizon 2020 research and innovation program under grant agreement No 951846.

** Early Stage Ph.D.

These approaches usually require the ontologies to fully expose their schema [2]. However, this might not always be feasible or fruitful. Our method is inspired from natural agent communication i.e. among animals, infants or humans, where the agents do not have access to each other's schema. Instead, these agents can only refer to each other's concepts by interacting with their environment and signaling each other. Language Games (LGs) are used by computational linguists to study and model how natural language emerges in order to cover communication needs among agents [13]. To that end, we are studying how LGs can be applied for Task-oriented Ontology Alignment (TOA), so that the agents can communicate using a newly invented language, without having to fully share their ontologies. Our method can be an alternative approach to TOA that does not require schema sharing, which, can even simplify the problem. This is the case when a layman and an expert interact, as for example, the interaction between a doctor and its patient, or a lawyer and its client. In such cases, although it might be feasible for them to exchange and align their ontologies in order to interact, this might be time-consuming, while also unnecessary depending on the task at hand. Of course, it is challenging to ensure that the new language is interpreted correctly by both parties, which is the focus of this thesis.

We suggest a LG approach, where the agents create a new vocabulary, the interpretation of which is aligned through interaction [14]. We propose such an approach because it has shown to allow successful communication between agents, without requiring them to expose any information explicitly. The interaction is designed in a way that requires the agents to interpret the same words with semantically similar properties across the schemata of different agents. Within LGs, agents are participating in a set of episodes in which they are required to act based on a sophisticated guess; a decision based on assumptions. The assumptions have the form of assertions, while the decision has an observable outcome which is either correct or wrong, i.e. the success of the episode. Depending on the episode's outcome, the agents either accept or reject their own assumptions, updating their knowledge and learn. We have completed early experiments applied on artificial cases reporting encouraging results that motivate us to extend our approach in order to apply it in real-world KGs.

2 State of the Art

In this section, we will introduce studies that are related to our work, ranging from the domain of LGs, where our suggested approach draws inspiration, to OA, where the approach can lead to useful outcomes.

2.1 Language Games

LGs are a common methodology for studying language emergence in populations of agents in a decentralised or self-organised way [13]. Language emerges through communicative interactions, allowing the agents to successfully and efficiently communicate while performing a specific task that the game defines [8]. Although the studies are applied on a set of agents, all interactions happen within communicative interaction rounds that involve two agents acting as

Speaker and Listener. The Speaker needs to communicate a piece of information to the Listener by uttering a word, so that the latter can make an informed decision. The word has the same form for the two agents, but they do not necessarily attribute the same meaning to it. Only the Speaker is allowed to create a new word when necessary, while all agents start with an empty vocabulary. The agents behave according to their word interpretations while attempting to cooperate. The communication is deemed successful according to the outcome of the task. In the case of success, the agents increase the probability of reusing the same interpretation in the future, or decrease it otherwise. Eventually, the agents converge to having shared interpretations allowing them to communicate successfully in any episode.

Referential Language Games. Most LG studies are designed around the referential task, where the Speaker needs to inform the Listener regarding the identity of a target object among a set of objects. In the simple, non-grounded, version of the game [14], the agents have their own names for every object in the world and must align words with the names that each agent gives to the same object. The agents make assumptions in the form of weights relating object names with words, which are updated after each interaction following a Reinforcement Learning (RL) setup driven by the outcome of the episodes. In the grounded version of the naming game, the agents are embodied and perceive objects in the form of features that have sub-symbolic values i.e. RGB colour, height, width, etc. Agents use words not to directly refer to an object, but implicitly do so by referring to its characteristics. This way, when the interactions end, the agents use the same word to refer to set of characteristics with similar enough values. Agents relate different characteristic values to the same word, depending on their observational position, the different lighting conditions, etc. This proves that this form of communication allows of bridging such a gap, leading to our motivation of attempting to align concepts that are similar enough across ontologies, not requiring them to have the exact same meaning.

2.2 Ontology Alignment

Traditional OA methods follow a centralised approach [2], where systems are asked to provide their complete knowledge in order to find plausible alignments. This can be a problem in scenarios where agents need to keep their knowledge private, or cannot share it for any other reason. Additionally, such techniques assume that it is known beforehand i) which systems will interact and ii) what types of tasks they will be asked to collaboratively perform. These are usually too strong assumptions, since different methods are resulting in different alignments and not all alignments are equally suitable for all tasks [12, 9]. Accordingly, two main streams of work have decided to propose techniques that allow the agents to develop alignments in a decentralised way through iterative agent interaction.

Negotiating Symbol Correspondences while Satisfying Private Constraints. This stream of studies suggests that agents interact in a set of episodes, during which they partly reveal their own knowledge and collaboratively converge to a common set of symbol correspondences [6, 7, 10, 11, 15]. These works formally define an argument, which consists of a suggested symbol correspondence together with supporting facts, as well as methods to generate and resolve them, even including rebuttals. The supporting facts either include previously accepted symbol correspondences or exposed facts from the agent's ontology. Each agent independently calculates a level of agreement with each suggested symbol correspondence, according to its ontology and the previously suggested correspondences, based on which, it decides whether to accept the suggestion or not. The method terminates when the agents cannot come up with new arguments and the commonly accepted symbol correspondences are provided as a solution. All these works show that even with partial revealings of the agents' ontologies, their method can achieve up to 95% of aligning accuracy [10], compared to centralised OA-based methods, where the complete information of both agents is accessible, making the task less challenging.

Symbol Correspondence Rectification via Agent Interaction. In a different stream of studies [1, 3–5], a population of agents engage in interaction rounds, called episodes, in order to repair or create a public set of symbol correspondences across their personal schemata. These studies are similar to our work because a population of agents engage in pairwise interactions describing an object and learn which symbols each uses for every property. Compared to our work, these studies expose their schemas and cannot create complex alignments. The experiments are performed in artificial ontologies that share partial information and do not contain any contradicting facts. Last but not least, these studies not only measure the evolution of the rate of successful communication as a communication criterion, but further evaluate the consistency, redundancy and other semantic measures, while also compare the produced alignments with a set of reference alignments in the form of recall and precision.

2.3 Differences with this Thesis

Compared to the presented OA studies, our LG-inspired TOA approach does not require the agents to expose any knowledge from their ontologies regarding their schema. Additionally, current studies are restricted to only produce simple alignments, not allowing their application on ontologies that are designed in different granularity, as is the interaction between an expert and a layman. For example, if one ontology defines a class Human for what the other ontology defines either as Woman or Man, simple alignments would face difficulties aligning these concepts. LG studies focus on studying language evolution among agents that sense the environment in an equally expressive way, e.g. centimeters and inches. Our aim is not to study language evolution, but to apply it as a method to perform a particular problem, namely Task-oriented Ontology Alignment (TOA). In such an application the ontologies are not expected to have the same expressiveness and communication success is not guaranteed. Finally, the agents in our case interpret words in terms of ontology concepts, the dependency of which must be taken into account. A summary of the comparisons is presented in Table 1.

Table 1. Comparing the suggested LG-inspired TOA approach to related work.

Method	Task	Assumptions	Exposed Knowledge
Negotiation [6]	Simple OA	Heuristic similarity values across ontologies	Minimised number of object properties
Rectification [1]	Simple OA	Sub-sampled ontologies from common ontology	Object properties
Referential LGs [13]	Study language Emergence	Common Knowledge	None
LG-inspired TOA	Complex TOA	Different Ontologies	Indirectly (by inference)

3 Problem Statement and Contributions

As presented earlier, we suggest that there is an overlapping interest between OA and LGs which has not been studied yet, leading us to our research question: "*Can Language Games be used for Task-oriented Ontology Alignment?*". We break our research question to four smaller ones:

1. To what extend can a LG approach perform TOA, without requiring the agents to reveal any of their schema?
2. What is the efficiency penalty or benefit imposed by restricting agents from exposing their knowledge?
3. Can such an approach be applied to multiple ontologies at the same time?
4. Can the agents extend their ontology or knowledge appropriately to always ensure successful communication?

Our contribution is a novel TOA method inspired by LGs, which will:

1. be able to deal with cases were communication success is not guaranteed;
2. extend current LG approaches to take into account the ontology's concept relationships as interpretation restrictions.
3. provide ontology alignments that are specifically tailored for a particular downstream task
4. broaden the application of OA methods to include cases where knowledge sharing is not possible, that can also be applied to ontologies of different granularity while even between more than two ontologies.

4 Research Methodology and Approach

Our method aims to indirectly align terms across ontologies, provided that these ontologies include information regarding a common set of instances identified as Uniform Resource Identifiers (URI), i.e. the **Common URIs**, as depicted at the center top of Figure 1. We assume this information to be in a Knowledge Graph (KG) form and more specifically in a Object - Property - Value triple format. We define a **characteristic** to be a property - value pair of a triple, so that the KG of an agent consists of objects and their characteristics. In Figure 1, characteristics are depicted as coloured property - value pairs, while same colours indicate

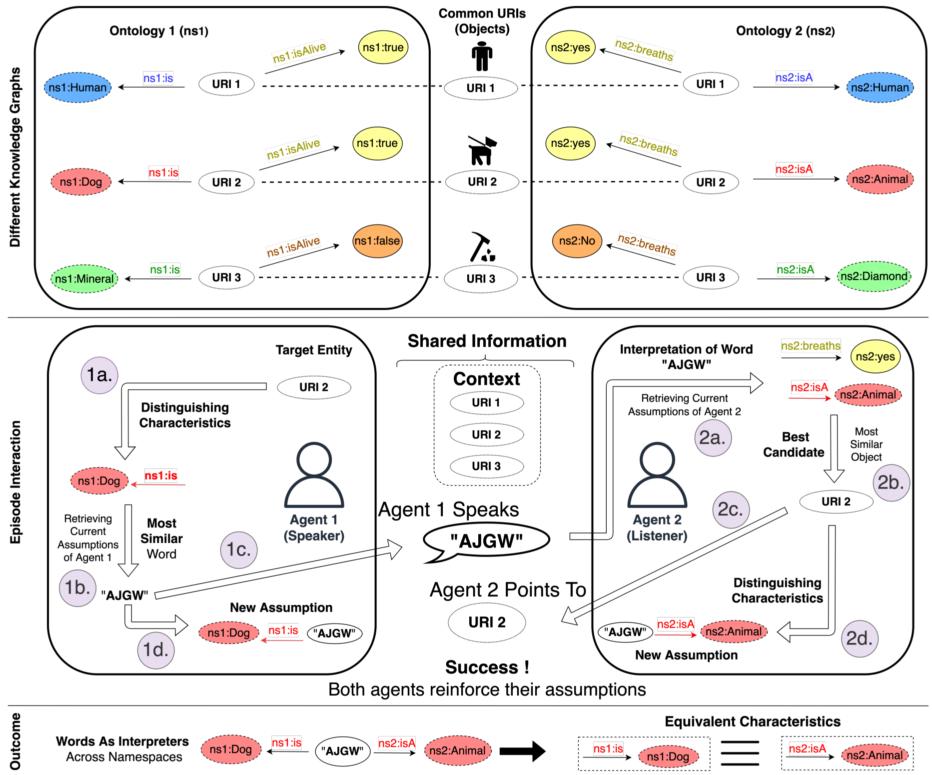

Fig. 1. (Top) Populated ontologies of Speaker and Listener. (Middle) Episode interaction; from top left to bottom right. (Bottom:) The alignment that this successful episode reinforces. The purple circles denote the methodology steps, while the colours indicate known characteristic alignments across ontologies.

ground truth equivalent characteristics across the two ontologies. Following the referential game setup and using the example from the middle Figure 1, Agent 1 acts as Speaker and needs to inform the Listener, impersonated by Agent 2, that the target object is URI 2 among the context objects i.e. URI 1, URI 2 and URI 3. Furthermore, we define the "**distinguishing characteristics**" of one object in the context as the subset of its characteristics that are not shared by any other object in the context, according to an agent's knowledge. Thus, they depend both on the context and the agent. On the middle left side of Figure 1 we can see the distinguishing characteristics of URI 2 i.e. {(ns1:is,ns1:Dog)}, according to Agent 1, with respect to the episode's context URIs. Within each episode the Speaker communicates one word and the Listener is allowed to communicate one URI in order to "point to" a candidate target item. The process is illustrated in the middle of Figure 1 starting from the Speaker knowing the target entity i.e. URI 2, and follows the arrow flow. Words are interpreted by relating them with characteristics and interpretations are different per agent. Both the objects and words are related in a boolean manner to a set of characteristics that are defined on the agent's schema. Accordingly, we define the **similarity** between words, objects or any set of characteristics, to be equal to the number of the characteristics that both are related to. We define every word-to-characteristic relation as a single **assumption**; see examples in Figure 1. When the agents converge to always communicating successfully, they can use the words to refer to each other's equivalent characteristics. At the bottom of Figure 1, we can see an example of the word "AJGW" being used as an interpreter across namespaces to suggest that its interpretations by the two agents, i.e. {(ns1:is,ns1:Dog)} and {(ns2:isA,ns2:Animal)}, are equivalent.

Episodes: Every agent can operate both as a Speaker or a Listener. An illustration of the episode interaction where Agent 1 and Agent 2 are acting as Speaker and Listener respectively, is given in the middle of Figure 1. The center depicts the information that is provided to both agents, while the sides are visualising the independent processes of the two agents that solely depend on their private information. As stated before, within each episode the Speaker needs to inform the Listener regarding the identity of the target object from a set of context objects. Every aspect of the episode is randomly sampled, i.e. the agents, their roles, the context and the target.

Speaker Behaviour: The Speaker initially calculates the distinguishing characteristics of the target object: which is {(ns1:is,ns1:Dog)} in the presented example (step 1a.). The same agent then uses its current assumptions to retrieve all words that are related to these characteristics resulting to a set of candidate words to communicate (step 1b.). All words that are more similar to any non-target object in the context are removed from this set, since communicating them would lead to misleading communication. In case the remaining set of words is empty, the speaker generates a new word, otherwise we select the most similar word from that set. Then, the Speaker communicates the selected word (step 1c.), and generates a set of assumptions relating the selected word with the distinguishing characteristics of the target object (step 1d.); i.e. one assumption per characteristic e.g."AJGW" $\rightarrow$ (ns1:is,ns1:Dog).

Listener Behaviour: The Listener interprets the communicated word by retrieving from memory all assumptions that relate characteristics to this word. In Figure 1, you can see the example of interpreting the word "AJGW" as the set of characteristics {(ns2:breaths,ns2:Dog), (ns2:isA,ns2:Animal) }; (step 2a). Note that each agent interprets a word in its own namespace. Similarity scores with context objects are calculated (step 2b.) while ties or lack of similarity scores, due to new words, are resolved randomly. The Listener points to the best matching context object (step 2c.) and generates a set of assumptions relating the communicated word with the distinguishing characteristics of the selected object (step 2d.); e.g. "AJGW" $\rightarrow$ (ns2:isA,ns2:Animal).

Outcome: The Speaker informs the Listener whether the object selection is correct. In case of successful communication, the two agents save the generated assumptions in their memories. Otherwise, they make sure that these assumptions are not in their memory. This way, the agents can use the interaction and the outcome of an episode in order to learn and update their interpretations. Eventually, the agents end up communicating successfully for enough subsequent episodes, at which point we can safely assume that the agents interpret the words using similar enough terms. For example, the interpretation assumptions that were generated during the depicted example episode interaction in Figure 1, allow the agents to refer to each other's equivalent characteristics $\{(\text{ns1:is,ns1:Dog})\} \equiv \{(\text{ns2:isA,ns2:Animal})\}$ by uttering the word "AJGW", as shown in the bottom of the figure. It should be mentioned that the agents to forget a word when it has not been used for the last 100 episodes, as it happens in the naming game.

5 Evaluation Plan

Proof of Concept. We run some proof of concept experiments towards answering our first research question as defined in Section 3. These investigate the successful application of our approach on small artificial ontologies. This stage has been performed and the results are presented in the next section.

RQ1. Next, we aim to apply our methodology to real ontologies, using existing TOA's benchmarks, and study how we need to improve our methodology appropriately. A candidate benchmark will either be found in the Ontology Alignment Evaluation Initiative, or will be constructed. We are aware that some episodes are more informative than others and we aim to allow the agents to design them, according to their state and goals. We estimate this phase to last around 6 months, helping us to answer our first research question.

RQ2. Towards answering out second research question, we will compare our method with other approaches that partly expose ontologies' knowledge and measure the amount of exposed information and the computational cost of each. This process should be performed within a period of 6 months.

RQ3. In order to apply our methodology to a population of ontologies, it is expected that the agents will need to have some theory of mind capabilities. This way, the agents will be able to learn what concepts they can communicate with every other agent separately. The experiments would be performed using benchmarks with more than 2 ontologies, and would help us answer our third research question i.e. the agent population setup. Estimated duration of this phase is around 9 months.

RQ4. The last phase will focus on assisting agents to decide when and how to extend their ontologies so that they learn each other's concepts, attempting to overcome ontology mismatch communication limitations. Experiments will be performed on the same benchmarks as before, but the agents should be able to communicate more concepts than before, if not all. This last phase should be performed within 1 year.

6 Preliminary Results

In this section, we will describe our proof of concept experiments. In each experiment, we provide all agents the same graph defined under different namespaces, except for the objects which have the same URIs across all ontologies. Thus, we know the ground truth alignments of the characteristics across graphs, allowing us to evaluate the output of our method beyond the success of the task. We run different experiments that vary according the number of object in the context (2 or 3) and the agent population size (2 or 3) depicted as "C. Size" and "Agents" respectively in the legends of Figure 2. Additionally, we use two small artificial KGs consisting of 10 and 20 triples describing 3 and 6 objects, denoted as "tiny" and "small".

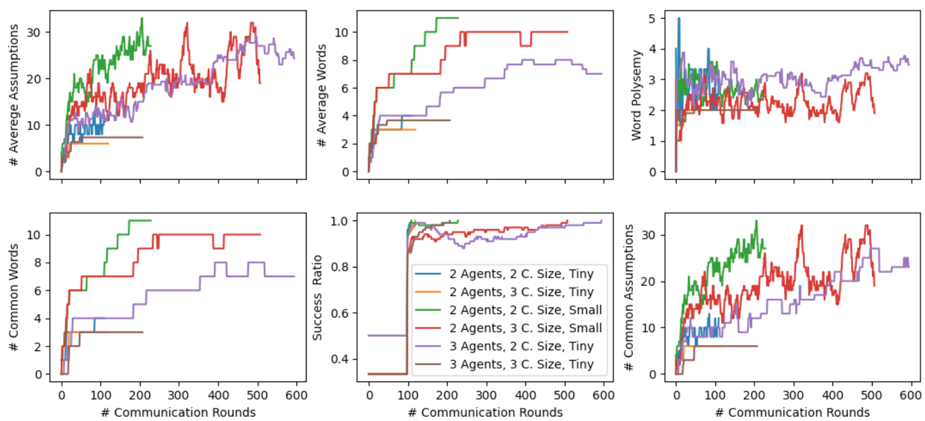

Fig. 2. The measured progression of the executed experiments, as captured by the 6 suggested evaluation metrics over the number of executed episodes.

Evaluation Metrics. We plan to monitor the progress of the experiments, over the number of episodes, in order to validate whether the agents converge to similar word interpretations and successful communication. Figure 2 presents the evolution of the suggested metrics over the number of completed episodes. We mainly want to measure the success regarding the task at hand. In the presented experiments the task is for the Listener to find the target object. The evaluation metric "Success Ratio" captures this, by calculating the average success value of the last 100 rounds. In case of less executed rounds, the metric is set to be equal to the chance of randomly selecting the target object: 1 / context size. Given that the agents are provided the same graphs, we can additionally measure the number current assumptions per agent and the number of common assumptions across agents, axes "#Average Assumptions" and "#Common Assumptions" respectively. We also present the traditional LG evaluation metrics to observe the successful application of the method. These are the average number of words ("#Average Words"), the number of common words across entities ("#Common Words") and the average number of characteristics related to one word ("Word Polysemy").

Table 2. An example of a converged experiment of how the communication symbols can be used as interpreters to align characteristics across namespaces (ns1 and ns2). Colours indicate ground truth characteristic correspondences across namespaces.

ns1 Characteristics		Word	ns2 Characteristics	
Property	**Value**		**Property**	**Value**
ns1:is	ns1:Human	"**DFHD**"	ns2:isA	ns2:Human
ns1:is	ns1:Dog	"**SEWG**"	ns2:isA	ns2:Animal
ns1:is	ns1:Dog	"**AJGW**"	ns2:isA	ns2:Animal
ns1:is ns1:isAlive	ns1:Mineral ns1:false	"**LWNC**"	ns2:isA ns2:breaths	ns2:Diamond ns2:No

The behaviour of the experiments according to our evaluation metrics are presented in Figure 2. First, it is important to note that all experiments converge to successful communication. This allows us to assume that our proof of concepts experiments were successful and motivates us to continue working on our approach. Furthermore, it seems that larger context sizes usually lead to faster convergence. This is intuitive since the larger number of objects within a context, the less number of distinguishing characteristics each object is expected to have, leading to less assumptions per communicated word. This leads to less average assumptions per agent which is also observed in the "#Average Assumptions" plot (orange, brown and red lines in Figure 2). It is important to point out that the methodology includes some stochasticity, and a more proper evaluation of the experiments would require aggregation over multiple executions of the same experiment. Regarding the alignments that the agents have generated, Table 2 shows how each agent interprets a word. The agents interpret the same words with the corresponding characteristics, even forming complex alignments (e.g. word "LWNC" on Table 2). On the other hand, the current methodology may generate synonyms i.e. two words with the same interpretation, which is an unwanted property (e.g. words "SEWG" and "AJGW" in Table 2).

7 Conclusions and Lessons Learned

To conclude, the success of the proof of concepts experiments suggests the continuation of our study, while also provide us with evidence of shortcomings that should be resolved. Specifically, the current method should be adjusted in order to avoid generating synonyms and the number of required episodes might be disproportional to the complexity of the problem, pointing towards further investigation for improvement. Therefore, the boolean relations between words and characteristics should be replaced to have a probabilistic form, defining a convex continuous space on which words are interpreted, as this allows easier optimisation. Furthermore, following the studies on OA, we should allow agents to store in their memory previous interactions, as a form of episodic memory, also aiming for faster convergence.

Acknowledgments. This work was performed in collaboration with dr. Ilaria Tiddi, prof. Annette ten Teije and prof. Frank van Harmelen. I want to sincerely thank them for their supervision and support.

References

1. van den Berg, L., Atencia, M., Euzenat, J.: A logical model for the ontology alignment repair game. Autonomous Agents and Multi-Agent Systems **35**, 32 (2021)
2. Euzenat, J., Shvaiko, P., et al.: Ontology matching, vol. 18. Springer (2007)
3. Euzenat, J.: First Experiments in Cultural Alignment Repair (Extended Version). In: Presutti, V., Blomqvist, E., Troncy, R., Sack, H., Papadakis, I., Tordai, A. (eds.) The Semantic Web: ESWC 2014 Satellite Events, vol. 8798, pp. 115–130. Springer International Publishing (2014)
4. Euzenat, J.: Crafting Ontology Alignments from Scratch Through Agent Communication. In: An, B., Bazzan, A., Leite, J., Villata, S., van der Torre, L. (eds.) PRIMA 2017: Principles and Practice of Multi-Agent Systems, vol. 10621, pp. 245–262. Springer International Publishing (2017), series Title: Lecture Notes in Computer Science
5. Euzenat, J.: Interaction-based ontology alignment repair with expansion and relaxation. In: Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence. pp. 185–191. International Joint Conferences on Artificial Intelligence Organization (2017)
6. Laera, L., Blacoe, I., Tamma, V., Payne, T., Euzenat, J., Bench-Capon, T.: Argumentation over ontology correspondences in MAS. In: Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems - AAMAS '07. p. 1. ACM Press, Honolulu, Hawaii (2007)
7. Laera, L., Tamma, V., Euzenat, J., Bench-Capon, T., Payne, T.: Arguing over ontology alignments. In: Proc. 1st ISWC 2006 international workshop on ontology matching (OM). pp. 49–60. Proc. 1st ISWC 2006 international workshop on ontology matching (OM) (2006)
8. Nevens, J., Eecke, P.V., Beuls, K.: A practical guide to studying emergent communication through grounded language games. CoRR **abs/2004.09218** (2020)
9. Payne, T.R., Tamma, V.: Negotiating over ontological correspondences with asymmetric and incomplete knowledge. In: Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems. pp. 517–524 (2014)
10. Payne, T.R., Tamma, V.: Using Preferences in Negotiations over Ontological Correspondences. In: Chen, Q., Torroni, P., Villata, S., Hsu, J., Omicini, A. (eds.) PRIMA 2015: Principles and Practice of Multi-Agent Systems, vol. 9387, pp. 319–334. Springer International Publishing (2015)
11. Santos, G., Tamma, V., Payne, T.R., Grasso, F.: Dialogue based meaning negotiation. In: The 15th Workshop on Computational Models of Natural Argument (CMNA 2015) (2015)
12. Shvaiko, P., Euzenat, J.: Ontology matching: State of the art and future challenges. Knowledge and Data Engineering, IEEE Transactions on **25**, 158–176 (2013)
13. Steels, L.: The emergence of grammar in communicating autonomous robotic agents. pp. 764–769 (01 2000)
14. Steels, L., Loetzsch, M.: The Grounded Naming Game pp. 41–59 (2012)
15. Trojahn, C., Euzenat, J., Tamma, V., Payne, T.R.: Argumentation for Reconciling Agent Ontologies. In: Kacprzyk, J., Elçi, A., Koné, M.T., Orgun, M.A. (eds.) Semantic Agent Systems, vol. 344, pp. 89–111. Springer Berlin Heidelberg (2011)