First experiments in cultural alignment repair (extended version)

Jérôme Euzenat

► To cite this version:

Jérôme Euzenat. First experiments in cultural alignment repair (extended version). 3rd ESWC workshop on Debugging ontologies and ontology mappings (WoDOOM), May 2014, Hersounisos, Greece. pp.115-130, 10.1007/978-3-319-11955-7_10 . hal-01180916

HAL Id: hal-01180916

<https://hal.science/hal-01180916>

Submitted on 8 Aug 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

First experiments in cultural alignment repair (extended version)*

Jérôme Euzenat

INRIA & LIG,
Grenoble, France
Jerome.Euzenat@inria.fr
<http://exmo.inria.fr>

Abstract. Alignments between ontologies may be established through agents holding such ontologies attempting at communicating and taking appropriate action when communication fails. This approach, that we call cultural repair, has the advantage of not assuming that everything should be set correctly before trying to communicate and of being able to overcome failures. We test here the adaptation of this approach to alignment repair, i.e., the improvement of incorrect alignments. For that purpose, we perform a series of experiments in which agents react to mistakes in alignments. The agents only know about their ontologies and alignments with others and they act in a fully decentralised way. We show that cultural repair is able to converge towards successful communication through improving the objective correctness of alignments. The obtained results are on par with a baseline of a priori alignment repair algorithms.

Keywords: Ontology alignment; alignment repair; cultural knowledge evolution; agent simulation; coherence; network of ontologies

1 Motivation

The work on cultural evolution applies, an idealised version of, the theory of evolution to culture. Culture is taken here as an intellectual artifact shared among a society. Cultural evolution experiments typically observe a society of agents evolving their culture through a precisely defined protocol. They perform repeatedly and randomly a task, called game, and their evolution is monitored. This protocol aims to experimentally discover the common state that agents may reach and its features. Luc Steels and colleagues have applied it convincingly to the particular artifact of natural language [11].

We aim at applying it to knowledge representation and at investigating some of its properties. A general motivation for this is that it is a plausible model of knowledge transmission. In ontology matching, it would help overcoming the limitations of current ontology matchers by having alignments evolving through their use [2], increasing the robustness of alignments by making them evolve if the environment evolves.

In this paper, we report our very first experiments in that direction. They consider alignments between ontologies as a cultural artifact shared among agents. In the following, we call cultural repair the action of agents repairing these alignments while trying to

* This paper is an extension of [6] with results of complementary experiments.

communicate. We hypothesise that it is possible to perform meaningful ontology repair with agents acting locally. The experiments reported here aims at showing that, starting from a random set of ontology alignments, agents can, through a very simple and distributed mechanism, reach a state where (a) communication is always successful, (b) alignments are coherent, and (c) their quality, as measured by F-measure, has increased. We also compare the obtained result to those of state-of-the-art repair systems.

Related experiments have been made on emerging semantics (semantic gossiping [4,3]). They involve tracking the communication path and the involved correspondences. By contrast, we use only minimal games with no global knowledge and no knowledge of alignment consistency and coherence from the agents. Our goal is to investigate how agents with relatively little common knowledge (here instances and the interface to their ontologies) can manage to revise networks of ontologies and at what quality.

2 Experimental framework

We present the experimental framework that is used in this paper. Its features have been driven by the wish that experiments be easily reproducible and as simple as possible. We first illustrate the proposed experiment through a simple example (§2.1), before defining precisely the experimental framework (§2.2) following [11].

2.1 Example

Consider an environment populated by objects characterised by three boolean features: $\text{color}=\{\text{white}|\text{black}\}$, $\text{shape}=\{\text{triangle}|\text{square}\}$ and $\text{size}=\{\text{small}|\text{large}\}$. This characterises $2^3 = 8$ types of individuals: $\blacksquare, \blacktriangle, \square, \triangle, \blacksquare, \blacktriangle, \square, \triangle$.

Three agents have their own ontology of what is in the environment. These ontologies, shown in Figure 1, identify the objects partially based on two of these features. Here they are a circular permutation of features: FC (shape, color), CS (color, size) and SF (size, shape).

In addition to their ontologies, agents have access to a set of shared alignments. These alignments comprise equivalence correspondences between their top (all) classes and other correspondences. Initially, these are randomly generated equivalence correspondences. For instance, they may contain the (incorrect) correspondence: $SF:\text{small} \equiv CS:\text{black}$.

Agents play a very simple game: a pair of agents a and b is randomly drawn as well as an object of the environment o . Agent a asks agent b the class c (source) to which the object o belongs, then it uses an alignment to establish to which class c' (target) this corresponds in its own ontology. Depending on the respective relation between c and c' , a may take the decision to change the alignment.

For instance, if agent CS draws the small-black-triangle ($\blacktriangle$) and asks agent SF for its class, this one will answer: small-triangle. The correspondence $SF:\text{small} \equiv CS:\text{black}$ and the class of $\blacktriangle$ in CS is black-small which is a subclass of $CS:\text{black}$, the result is then a SUCCESS. The fact that the correspondence is not valid is not known

Fig. 1. Example of a generated network of ontologies with the exact reference alignments.

to the agents, the only thing that counts is that the result is compatible with their own knowledge.

If, on the contrary, the drawn instance is small-white-triangle ($\triangle$), *SF* would have made the same answer. This time, the result would be a FAILURE because $\triangle$ belongs to class *CS*:white-small which is disjoint from *CS*:black-small.

How to deal with this failure is a matter of strategy:

- delete** *SF*:small $\equiv$ *CS*:black can be suppressed from the alignment;
- replace** *SF*:small $\equiv$ *CS*:black can be replaced by *SF*:small $\leq$ *CS*:black;
- add** in addition, the weaker correspondence *SF*:small $\geq$ *CS*:all can be added to the alignment (but this correspondence is subsumed by *SF*:all $\equiv$ *CS*:all).

In the end, it is expected that the shared alignments will improve and that communication will be increasingly successful over time. Successful communication can be observed directly. Alignment quality may be assessed through other indicators: Figure 1

shows (in dotted lines) the correct (or reference) alignments. Reference alignments are not known to the agents but can be automatically generated and used for measuring the quality of the resulting network of ontologies by comparison with it.

2.2 Experimental set up

We systematically describe the different aspects of the carried out experiments in the style of [11].

Environment: The environment contains objects which are described by a set of n characteristics (we consider them ordered). Each characteristic can take two possible values which, in this experiment, are considered exclusive.

Population: The experiment uses n agents with as many ontologies. Each agent is assigned one different ontology. In this first setting, each agent will have an ontology based on $n - 1$ of these characteristics (each agent will use the first $n - 1$ characteristics starting at the agent's rank). The ontology is a simple decision trees of size 2^{n-1} in which each level corresponds to a characteristic and subclasses are disjoint.

Shared network of ontologies: A complete network of $\frac{n \times (n-1)}{2}$ invertible alignments between the ontologies is shared among agents (public). The network is symmetric (the alignment between o and o' is the converse of the alignment between o' and o) and a class is in at most one correspondence per alignment.

Initialisation: In the initial state, each alignment contains equivalence correspondences between the most general classes of both ontologies, plus 2^{n-1} randomly generated equivalence ($\equiv$) correspondences.

Game: A pair of distinct agents $\langle a, b \rangle$ is randomly picked up as well as a set of characteristic values describing an individual (equiprobable). The first agent (a) asks the second one (b) the (most specific) class of its ontology to which the instance belongs (*source*). It uses the alignment between their respective ontologies for finding to which class this corresponds in its own ontology (*target*). This class is compared to the one the instance belongs to in the agent a ontology (*local*).

Success: Full success is obtained if the two classes (*target* and *local*) are the same. But there are other cases of success:

- *target* is a super-class of *local*: this is considered successful (this only means that the sets of alignments/ontologies are not precise enough);
- *target* is a sub-class of *local*: this is not possible here because for each instance, *local* will be a leaf.

Failure: Failure happens if the two classes are disjoint. In such a case, the agent a will proceed to repair.

Repair: Several types of actions (called modalities) may be undertaken in case of failure:

delete the correspondence is simply discarded from the alignment;

replace if the correspondence is an $\equiv$ correspondence it is replaced by the $\leq$ correspondence from the target class to the source class;

add in addition to the former a new $\leq$ correspondence from the source to a superclass of the target is added. This correspondence was entailed by the initial correspondence, but would not entail the failure.

These three strategies share two properties: (a) the resulting alignment would have avoided the uncovered mistake; (b) the resulting alignment is logically entailed by the initial one.

Success measure: The classical success measure is the rate of successful communication, i.e., communication without failure.

Secondary success measure: Several measures may be used for evaluating the quality of the reached state: consistency, redundancy, discriminability. We use two different measures: the averaged degree of incoherence [9] and the semantic F-measure [5]. Indeed, this setting allows for computing automatically the reference alignment in the network, so we can compute F-measure.

External validation: The obtained result can be compared with that of other repair strategies. We compare the results obtained with those of two directly available repair algorithms: Alcomo [8] and LogMap repair [7].

3 Experiments

We report four series of experiments designed to illustrate how such techniques may work and what are their capabilities. In addition to these experiments, which were reported in [6], we have added extra confirmatory experiments which test hypotheses stemming from the observations.

The tests are initially carried out on societies of at least 4 agents because, in the setting with 3 agents, the delete modality drives the convergence towards trivial alignments (containing only $\text{all} \equiv \text{all}$) and the other modalities do it too often.

All experiments have been run in a dedicated framework that is available from <http://lazylav.gforge.inria.fr>.

3.1 Convergence

We first test that, in spite of mostly random modalities (random initial alignments, random agents and random instances in each games), the experiments converge towards a uniform success rate.

Four agents are used and the experiment is run 10 times over 2000 games. The evolution of the success rate is compared.

3.2 Modality comparison

The second experiment tests the behaviour of the three repair modalities: delete, replace, add.

Four agents are used and the experiment is run 10 times over 2000 games with each modalities. The results are collected in terms of average success rate and F-measure.

3.3 Baseline comparison

Then the results obtained by the best of these modalities are compared to baseline repairing algorithms in terms of F-measures, coherence and number of correspondences.

The baseline algorithms are Alcomo and LogMap repair. The comparison is made on the basis of success rate, F-measure and the number of correspondences.

LogMap and Alcomo are only taken as a baseline: on the one hand, such algorithms do not have the information that agents may use, on the other hand, agents have no global view of the ontologies and knowledge of consistency or coherence.

3.4 Scale dimension

Finally we observe settings of increasing difficulty by taking the modality providing the best F-measure and applying it to settings with 3, 4, 5 and 6 ontologies.

This still uses 10 runs with the add modality over 10000 games. Results are reported as number of correspondences, F-measure and success rate and compared with the best F-measure of Alcomo and LogMap.

4 Results

Results of the four presented experiments are reported and discussed.

4.1 Convergence

Figure 2 shows the result of our first experiment: 10 runs with a random network as defined above with 4 ontologies. Each curve corresponds to one of the 10 runs over 2000 iterations.

Fig. 2. Ten random runs and their overall success rate, i.e., the proportion of games which were successful so far [mod=add; #agents=4; #games=2000; #runs=1].

Figure 2 shows a remarkable convergence between the runs. After the first 200 games dominated by randomness, they converge asymptotically and at the same pace towards 100%. Indeed, as soon as the network of ontologies has been cleaned up (around 1200 iterations maximum), the rate only grows. It never reaches 1 because of the initial period which contains failures.

From now on, we will still consider 10 runs, but the results will be averaged over these runs.

4.2 Modality comparison

Figure 3 shows the evolution over 2000 iterations of the success rate and F-measure of the three presented modalities.

Fig. 3. The average F-measures (dashed) and success rate (plain) with the three different modalities: delete (red), replace (green) and add (blue) [mod=del,repl,add; #agents=4; #games=2000; #runs=10].

delete converges more quickly than replace which converges more quickly than add. This can easily be explained: delete suppresses a cause of problem, replace only suppresses half of it so it may need one further deletion for converging, while add replaces one incorrect correspondence by two correspondences which may be incorrect, so it requires more time to converge.

For the same reason, the success rate is consequently higher. Table 1 shows that for the delete modality, 97.6% success rate corresponds to 48 failure, i.e. 48 deleted correspondences over 54. The 6 remaining correspondences are all $\equiv$ all correspondences. replace reaches the same result with a 95.2% rate, which corresponds to twice as many failures.

The results of delete and replace modalities are the same: in order to be correct, alignments are reduced to the all $\equiv$ all correspondences. This is unavoidable for delete (because initial correspondences are equivalences, although, by construction, the correct correspondences are subsumption, so the initial correspondences are incorrect in at least one direction). This is by chance, and because of averaging, for replace.

On the contrary, the add modality has a 88.6% success rate, i.e., 228 failures. This means that on average for each correspondence it has generated 4 alternative correspon-

dences. This is only an average because after 2000 games (and even after 10000 games), there remain more than 12 correspondences.

Contrary to the other modalities, add improves over the initial F-measure.

Table 1 shows that all methods reach full consistency (incoherence rate=0.) from a network of ontologies with 50% incoherence, i.e., half of the correspondences are involved in an inconsistency (or incoherence).

Modality	Size	Success		Semantic F-measure	Syntactic F-measure	Convergence
		rate	Incoherence degree			
reference	70	-	0.0	1.0	1.0	-
initial	54	-	[0.46-0.49]	0.20	(0.20)	-
delete	6	0.98	0.0	0.16	(0.16)	400
replace	6	0.95	0.0	0.16	(0.16)	1000
add	12.7	0.89	0.0	0.23	(0.16)	1330
Alcomo	25.5	-	0.0	0.26	(0.14)	-
LogMap	36.5	-	0.0	0.26	(0.14)	-

Table 1. Results of the three different modalities compared with Alcomo and LogMap on 10 runs, 4 ontologies and 2000 iterations. Syntactic F-measure has been obtained in an independent but identical evaluation.

Concerning F-measure, add converges towards a significantly higher value than the two other approaches. With four ontologies, it has a chance to find weaker but more correct correspondences. The add strategy is more costly but more effective than the two other strategies.

4.3 Baseline comparison

This experiment exploits the same data as the previous one (§4.2); exploiting those of the next experiment (on 10000 iterations) provides similar results.

Table 1 shows that all three methods are able to restore full coherence and to slightly improve the initial F-measure. Their result is overall comparable but, as can be seen in Figure 4, cultural repair does not reach the F-measure of logical algorithms.

Cultural repair finds half of the correspondences of Alcomo and one third of those of LogMap. This is expected because Alcomo only discards the minimum number of correspondences which bring incoherence, while LogMap weakens them (like the add modality). The agents having more information on what is incorrect, discard more correspondences.

When looking at F-measures, it seems that logical repair strategies can find more than 6 new correspondences which are correct while the add strategy can only find more than 3. This is not true, as shown in Table 1, because we use *semantic* precision and recall [5]. Figure 5 shows that in terms of syntactic F-measure, the cultural revision approach does better than the logical baselines. These methods preserve correspondences which are not correct, but which entails correct correspondences. This increases semantic recall and F-measure.

Fig. 4. Average success, semantic F-measure and number of correspondences for the **add** modality compared to the Alcom and LogMap F-measure as a baseline [mod=add; #agents=4; #games=2000; #runs=10].

Fig. 5. Average success, syntactic F-measure and number of correspondences for the **add** modality compared to the Alcom and LogMap F-measure as a baseline [mod=add; #agents=4; #games=2000; #runs=10].

There is a large variation on the results given by the different methods. Out of the same 10 runs, LogMap had the best F-measures 5 times, Alcomo 3 times, and cultural repair twice. But the largest variation is obtained by cultural repair with a F-measure ranging from 0.16 to 0.33. Its result is indeed highly dependent on the initial alignment.

4.4 Scale dimension

So far, we concentrated on 4 agents, what happens with a different number of agents? The number of agents does not only determine the number of ontologies. It also determines the number of alignments (quadratic in the number of ontologies), the number of correspondences per alignments and the number of features per instances. This means that the more agents are used, the slower is the convergence. So, we played 10000 games in order to have a chance to reach a satisfying level of F-measure.

Figure 6 shows the regular pattern followed by cultural repair: the first phase is random and increases the number of correspondences (due to the *add* modality). Then, this number slowly decreases. Cultural repair is slower to converge as the problem size increases. This is easily explained: as the correction of the alignment converges, the number of failure-prone games diminishes. Since games are selected at random, the probability to pick up the last configurations (in the end there is only one) becomes lower and lower. The increased number of iterations to converge is directly tied to the largely increased difficulty of the task (number of agents, number of alignments, size of ontologies, characteristics of objects).

This increase is not a measure of the complexity of the approach itself. In fact, it is highly distributed, and it is supposed to be carried out while agents are achieving other tasks (trying to communicate). All the time spent between the last two failures is time of communicative *success*, i.e., agents never had to suffer from the wrong correspondences.

A very simple strategy for improving this would be that agents try to select themselves examples in order to verify the correspondences that they have not already tested.

# agents	# correpondences					Incoherence				F-measure				Convergence
	Reference	Initial	LogMap	Alcomo	Final	Initial	LogMap	Alcomo	Final	Initial	LogMap	Alcomo	Final	
3	15	15	12	10.3	3	0.31	0.	0.	0.	0.32	0.35	0.36	0.33	300
4	70	54	36.7	28.4	12.4	0.47	0.	0.	0.	0.20	0.24	0.25	0.21	1670
5	250	170	94.7	71.7	47.4	0.58	0.	0.	0.	0.11	0.18	0.17	0.24	5400
6	783	495	234	182	224	0.63	0.	0.	0.	0.06	0.12	0.11	0.14	10.000+

Table 2. Number of correspondences, incoherence rate and F-measure over 10000 games.

Fig. 6. 10.000 games with 3, 4, 5 and 6 ontologies [mod=add,#agents=3,4,5,6; #games=10000; #runs=10].

5 Additional experiments related to scaling

The first experiments are small games with 3 agents, ontologies with 7 classes and alignments with 5 correspondences. The results are instantaneous.

If we augment the number n of agents, we also have programmed for augmenting the number of classes ($n^2 - 1$), the number of alignments ($\frac{n \times (n-1)}{2}$) and the number of correspondences per alignment ($2^{n-1} + 1$). Then convergence is slower.

However, we may think of unleashing the relation between these quantities and the number of agents. It is also possible to augment the number of agents with agents sharing the same set of ontologies. This should lead to different convergence profiles.

Here we consider two specific extensions of the previous experiments focusing on their results in terms of F-measure. We also run an experiment reducing the number of disjointness statements in the ontologies, but this experiment was non conclusive.

5.1 Playing over 10000 games

Table 2 seems to show that, as the complexity of the problem increases, the F-measure of cultural repair is better than that of logical repair mechanisms. We have confirmed this by running the last experiment (with 6 agents) over 100000 games (Table 3) to be sure to reach convergence. In fact, convergence occurs quite early: out of 10 runs the last failure occurs at game 16957, but the 10 last failures occur between 11352 and 16957 (the second one is at game 13706).

This experiment shows that at convergence, F-measure is better than the baseline.

Modality	Size	Success		Semantic	
		rate	degree	F-measure	Convergence
reference	783	-	0.0	1.0	-
initial	495	-	0.64	0.06	-
add	208	0.97	0.0	0.15	17000
Alcomo	180	-	0.0	0.11	-
LogMap	227	-	0.0	0.12	-

Table 3. Results of the add modality with Alcomo and LogMap on 10 runs, 6 ontologies and 100000 iterations.

5.2 Augmenting the number of correspondences in the initial alignments

The initial number of correspondences per alignments ($2^{n-1} + 1$) had been chosen as an approximation of the number of correspondences in the reference alignment. As can be observed from Table 2, this approximation underestimates that number. We run additional experiments with the usual 4 agents to observe the effect of increased initial alignments.

From Figure 7, we observe as expected that the success rate is lower to converge (there are more correspondences to invalidate), but F-measure grows to a higher level (the increased number of correspondences increases the chances to retain a correct correspondence).

Fig. 7. Success rate and F-measure obtained with the **add** modality with varying number of initial correspondences in alignments [#coresp=6,8,10,11,12,mod=add,#agents=4; #games=2000; #runs=10].

Figure 8 shows that with an increased number of correspondences in initial alignments, the resulting F-measure is increased for cultural repair while it decreases in the baseline. The reason is that this increases the chances of the former to retain a correct correspondence, while this also increases the chances of the latter to retain an incorrect correspondence. Indeed, the former will in the end eliminate all incorrect correspondences while the latter tries to preserve as many correspondences as possible as long as the alignments are coherent.

6 Discussion

The relatively low F-measure rate is tied to the type of experiments: agents are constrained by the initial alignment. They do not invent any correspondences, they only repair them by reducing their scope. In this respect, they are on par with logical repair algorithms.

However, they have more information than these repair algorithms. It could then be expected that their results are higher. This is not the case because, when an initial correspondence is unrelated to the valid one, agents will simply discard them. They will thus end up with few correspondences with a high precision and low recall.

The state-of-the-art repair algorithms will preserve more correspondences because their only criterion is consistency and coherence: as soon as the alignment is coherent,

Fig. 8. Semantic F-measure obtained with the `add` modality with varying number of initial correspondences in alignments compared to LogMap and Alcomo [`#coresp=6,8,10,11,12,mod=add,#agents=4; #games=2000; #runs=10`].

such algorithms will stop. One could expect a lower precision, but not a higher recall since such algorithms are also tied to the initial alignment.

Because we use semantic precision and recall, it happens that among these erroneous correspondences, some of them entail some valid correspondences (and some invalid ones). This contributes to increase semantic recall.

7 Conclusion

We explored how mechanisms implemented as primitive cultural evolution can be applied to alignment repair. We measured:

- Converging success rate (towards 100% success);
- Coherent alignments (100% coherence);
- F-measures on par with logical repair systems;
- A number of games necessary to repair increasing very fast.

The advantage of this approach are:

- It is totally distributed: agents do not need to have the knowledge of what is an inconsistent or incoherent alignment (only an inconsistent ontology).
- The repair of the network of ontologies is not blind, i.e., restoring inconsistency without knowing if it is likely to be correct, so it also increases F-measure (which is not necessarily the case of other alignment repair strategies [10]).

Yet, this technique does not replace ontology matching nor alignment repair techniques.

8 Perspectives

We concentrated here on alignment repair. However, such a game can perfectly be adapted for matching (creating missing correspondences and revising them on the fly).

In the short term, we would like to adapt this technique in two directions:

- introducing probabilities and using such techniques in order to learn confidence on correspondences that may be used for reasoning [1],
- dealing with alignment composition by propagating instances across agents in the same perspective as the whispering games (propagating classes and see what comes back, setting weights to correspondences) [4].

In the longer term, such techniques do not have to be concentrated on one activity, such as alignment repair. Indeed, they are not problem solving techniques (solving the alignment repair problem). Instead, they are adaptive behaviours, not modifying anything as long as activities are carried out properly, and reacting to improper situations. So, cultural knowledge evolution techniques must be adapted to the lifetime of a society of agents which can generate alignments, react to environment change and encounter other societies.

9 Acknowledgements

Thanks to Christian Meilicke (Universität Mannheim) and Ernesto Jimenez Ruiz (Oxford university) for making Alcomo and LogMap available and usable. Thanks to an anonymous reviewer for useful suggestions.

References

1. Manuel Atencia, Alexander Borgida, Jérôme Euzenat, Chiara Ghidini, and Luciano Serafini. A formal semantics for weighted ontology mappings. In *Proc. 11th International Semantic Web Conference (ISWC)*, volume 7649 of *Lecture notes in computer science*, pages 17–33, Boston (MA US), 2012.
2. Manuel Atencia and Marco Schorlemmer. An interaction-based approach to semantic alignment. *Journal of Web Semantics*, 13(1):131–147, 2012.
3. Thomas Cerqueus, Sylvie Cazalens, and Philippe Lamarre. Gossiping correspondences to reduce semantic heterogeneity of unstructured P2P systems. In *Proc. 4th International Conference on Data Management in Grid and Peer-to-Peer Systems, Toulouse (FR)*, pages 37–48, 2011.
4. Philippe Cudré-Mauroux. *Emergent Semantics: Interoperability in large-scale decentralized information systems*. EPFL Press, Lausanne (CH), 2008.
5. Jérôme Euzenat. Semantic precision and recall for ontology alignment evaluation. In *Proc. 20th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 348–353, Hyderabad (IN), 2007.
6. Jérôme Euzenat. First experiments in cultural alignment repair. In *Proc. 3rd ESWC workshop on Debugging ontologies and ontology mappings (WoDOOM), Hersounisos (GR)*, pages 3–14, 2014.

7. Ernesto Jiménez-Ruiz, Christian Meilicke, Bernardo Cuenca Grau, and Ian Horrocks. Evaluating mapping repair systems with large biomedical ontologies. In *Proc. 26th Description logics workshop*, 2013.
8. Christian Meilicke. *Alignment incoherence in ontology matching*. PhD thesis, Universität Mannheim, 2011.
9. Christian Meilicke and Heiner Stuckenschmidt. Incoherence as a basis for measuring the quality of ontology mappings. In *Proceedings of the 3rd ISWC international workshop on Ontology Matching*, pages 1–12, 2008.
10. Catia Pesquita, Daniel Faria, Emanuel Santos, and Francisco Couto. To repair or not to repair: reconciling correctness and coherence in ontology reference alignments. In *Proc. 8th ISWC ontology matching workshop (OM), Sydney (AU)*, pages 13–24, 2013.
11. Luc Steels, editor. *Experiments in cultural language evolution*. John Benjamins, Amsterdam (NL), 2012.