

Crafting ontology alignments from scratch through agent communication

Jérôme Euzenat

Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG, F-38000 Grenoble France
<http://mOeX.inria.fr>

Abstract. Agents may use different ontologies for representing knowledge and take advantage of alignments between ontologies in order to communicate. Such alignments may be provided by dedicated algorithms, but their accuracy is far from satisfying. We already explored operators allowing agents to repair such alignments while using them for communicating. The question remained of the capability of agents to craft alignments from scratch in the same way. Here we explore the use of expanding repair operators for that purpose. When starting from empty alignments, agents fails to create them as they have nothing to repair. Hence, we introduce the capability for agents to risk adding new correspondences when no existing one is useful. We compare and discuss the results provided by this modality and show that, due to this generative capability, agents reach better results than without it in terms of the accuracy of their alignments. When starting with empty alignments, alignments reach the same quality level as when starting with random alignments, thus providing a reliable way for agents to build alignment from scratch through communication.

1 Introduction

We usually develop a representation of the world we live in from a combination of experience and education. Artificial agents may be designed along the same principles. Experience is gathered through the interaction with their environment and education through the interaction with their peers.

Various experiments have been developed concerning how interaction with peers can shape agents knowledge. In particular, we dealt with situations in which agents have correct but not complete description of their environments as ontologies and they use alignments between each others ontologies to interact. Alignments between two ontologies are made of correspondences which express the (subsumption) relation between different concepts of these ontologies.

For instance, agents may classify the objects of their environment through a collection of features: size, colour, ability to move, dangerousness, etc. They may however classify such objects differently: first by colour, then by dangerousness or first by mobility, then by size. In such conditions, it can be difficult to communicate. Agents may use alignments between their ontologies expressing that all purple dangerous objects are mobile or that static objects are yellow and harmless. The question we explore is how such alignments may be elaborated.

We have developed experiments in which agents can repair random alignments through playing some interaction game [7]. For that purpose, agents would use adaptation operators to alter the alignments when their use causes errors. For instance, discovering an object which is yellow, harmful, static and small would invalidate the second correspondence. An operator may thus discard it and replace it by a weaker correspondence such as static and large objects are yellow and harmless. However, the experimental protocol was relying on random initial alignments that agents would only correct.

Recently, such operators have been improved by providing agents with the opportunity to go beyond repairing existing alignments by introducing new tentative correspondences when necessary [8]. This was shown to improve results measured as the proximity of the resulting alignments with those known as correct. It remained unclear whether agents could build correspondences *ex nihilo* through communication. However, the newly introduced modalities open the possibility to test agents starting with no alignment and generating tentative (random) correspondences when they do not know any.

Here, we show how such a function can be implemented in artificial agents and observe the way such agents behave. We first consider what it means to start with no alignments for communicating and find that, in such a case, agents make no progress as they do not have an occasion to discover errors to correct. Hence, we introduce a new variation in which agents having no non-trivial correspondence to use would generate one tentatively useful one on the spot.

We compare results obtained with this new modality when agents are provided with random correspondences or not. It happens that, already if agents are provided with random initial correspondences, this new modality improves the results. Moreover, agents without alignments can develop them from scratch through interacting and reach results with the same quality as agents provided with random alignments (and the same modalities).

The paper is organised as follows. After reviewing related work (§2), we recall the settings previously developed for studying alignment evolution (§3). We then report on starting with empty networks of ontologies (§4), using the new generative modality (§5) and combining both approaches (§6).

2 Related work

Interacting agents need a way to understand each others to some degree. They can all use the same ontology or preserve heterogeneous ontologies. In the latter case, heterogeneity is often reduced by using alignments between these ontologies [13]. However, because matching systems are not always correct or because agents change their ontologies, alignments may become incorrect. There have been various ways to solve this problem independently of agent tasks: gossiping to ensure a global coherence of the networks of ontologies [1], arguing over correspondences to select the relevant ones [17, 13], logical repair to enforce consistency [11, 9, 14] or conservativity [10] constraints. Such approaches have been integrated within agent systems through specific protocols [1, 18, 13].

Cultural language evolution [16] showed how a particular culture can be shared by a population of agents through communication. This work offers an experimental methodology: a population of agents has to play randomly an “interaction game” with precise rules. The outcome of the game is clearly identified as a success or a failure. In function of the outcome, agents adapt their language. The state of the system is monitored and especially the success rate which measures the convergence of agents to a stable state. This approach may be applied to language [16] or ontologies [15].

In this paper, we are considering approaches in which agents elaborate and repair alignments, playing such interaction games, achieving global repair through local action.

This approach has been applied to agent-based data interlinking where agents would exchange graph patterns assumed to characterise entities to be identified [2]. They compared several ways to select the graph patterns to be exchanged in order to maximise the recognition. It has also been considered in the context of more elaborate games. Interaction-situated semantic alignment [3] considers ontology matching as framed by interaction protocols that agents use to communicate. Agents induce alignments between the different ontologies that they use depending on the success expectation of each correspondence with respect to the protocol. Failing dialogues lead them to revise their expectations and associated correspondences.

Cultural evolution has been adapted to the evolution of ontology alignments: agents have their own ontologies related by public ontology alignments (network of ontologies). They play simple reclassification games involving the alignments. The games allow agents to detect incorrect correspondences through their use and to locally repair the alignment with adaptation operators. Several operators were compared and were shown to converge towards fully correct alignments [7, 8]. The evolutionary approach has also been combined with classical matching, but can be interpreted as a way to repair alignments through their use [4]. This has recently been generalised to less constrained declarative protocols with isomorphic alignments [5].

Anemone [18] is a comprehensive framework with the same goal: matching ontologies opportunistically at run time. It provides a protocol allowing agents to negotiate the relation between concepts of their ontologies when they cannot express a message they know the interlocutor will understand. Agents can then communicate a definition of the concept or, if this is not possible, provide examples of concept instance.

In spite of superficial differences —we repair alignments, Anemone extends ontologies—, both approaches may be considered as achieving the same tasks, in a lazy way, and they share their inspiration from cultural evolution for the assertional aspect. However, there are also notable differences in the proposed concepts and the methodology. Regarding concepts, our agents never exchange concept definitions, nor rely on negotiation. They operate at a lower level in which they apply corrective operations based on simple failure detection. As a consequence, communication is sound but not lossless: agents usually end up in a more general concept than what can be expressed. Concerning methodology, some theoretical properties (sound and lossless communication) of Anemone have been established and the system has been illustrated on a use case [18]. Here, we provide randomised experimental results comparing various oper-

ators and modalities which may be considered as variants of the Anemone protocol. Hence these work may be seen as complementary.

The work considered here is an extension of [7, 8] to the case where agents starts with empty alignments. It can be thus considered as another instance of ontology matching through interaction [3, 4]. However, instead of being guided by the interaction protocol, it is guided by the more expressive ontology structure.

3 Experimental setting and background

The experimental method and software used in this work is directly derived from [7]. The setting is deliberately abstract. This allows to control experiment variables, as our goal is to understand better which factors, in the setting and adaptation operators, influence the properties of the result. These experiments are randomised, which would be difficult with real-world data.

3.1 Setting

Objects in the world are described by a finite set of Boolean features (named a , b , c , etc.). Each object is described by the presence or absence of each feature, e.g., $a \wedge \neg b \wedge \neg c \wedge d$.

Objects, also called instances, are classified in ontologies which are strict hierarchies. Each class in the hierarchy corresponds to the conjunction of the features of its ancestors. For instance, the bottom-leftmost class in Figure 1 is defined by $\neg c \wedge \neg d \wedge \neg a$. Each level in these hierarchies adds one constraint (positive or negative) related to exactly one feature. This means that classes not in subsumption relation are disjoint. Ontologies are incomplete since they have one level less than the environment has features (d for the middle ontology of Figure 1). Ontologies, expressed in OWL, only use a very simple description logic. The boolean separation, may seem universal and/or artificial, it is simply a minimal language that features subsumption and disjointness.

An alignment is a set of correspondences between two ontologies. A correspondence is an expression like $\langle C, r, C' \rangle$ in which C and C' identify classes of the two ontologies and r is the relation between these classes. We use relation symbols \leq , \geq and $=$ denoting subsumption and equivalence relations between classes. Figure 1 shows a fragment of the correct alignments, called reference, between one ontology and two others. Alignments are required to be functional—a class is subsumed by at most one more specific class in each alignment, but one class may subsume several classes—so as to be able to use them to reclassify instances.

The instance \blacktriangle featured in Figure 1 is characterised by features $\neg a \wedge b \wedge \neg c \wedge \neg d$. Knowing this, each agent can classify it in its own ontology under the adequate most specific class. By using the expressed alignments, agents with the right most ontology, can reclassify this instance in the class $\neg a \wedge b$ in the middle ontology.

3.2 Games

There are as many agents as ontologies (each agent has a different ontology). Agents do not have access to the ontologies of other agents—they know the class names but




Fig. 1. One central ontology and fragments of the correct alignments to two other ontologies.

not their definitions— nor to the reference alignments. They are instead provided with random alignments with the constraint that the topmost classes of each ontology are equivalent and alignments are functional.

Agents are no other goal than continuously playing an interaction game. In each game, two agents are chosen at random, one agent (A) picks up an instance (\blacktriangle) at random and asks the other one (A') in which class it would classify the instance in its own ontology (O) using the public alignment between O and O' . A' determines which correspondence is applicable and communicates both the correspondence ($\langle C, r, C' \rangle$) and the class (C). A considers the relations between C and the class D in which it would classify \blacktriangle .

If C is compatible with D (C subsumes D , noted $D \sqsubseteq C$) the interaction is considered successful, otherwise (C is disjoint from D , noted $C \perp D$) it is a failure. A and A' then use an adaptation operator to adapt the alignment (in this case, they only do it in case of failure). Various adaptation operators may be used.

3.3 Adaptation operators

The way agents play the game is built into their design. They use a specific adaptation operator to adapt their alignments. We consider six operators [8]: delete, replace, refine, add, addjoin and refadd. Assuming that the faulty correspondence $\langle C, r, C' \rangle$ has been crossed by the object from C' to C :

delete simply removes the correspondence;

replace in case r is $=$, then, in addition to delete, adds the same correspondence with a \leq relation ($\langle C, \leq, C' \rangle$);

refine extends replace by adding a correspondence between C and the subclasses C'' of C' that do not subsume the actual class of the object ($\langle C, \geq, C'' \rangle$);

add extends replace by adding a correspondence between C' and the direct superclass C'' of C ($\langle C'', \geq, C' \rangle$);

addjoin is a variation of add which adds a correspondence between C' and the lowest superclass C'' of C compatible with D ($\langle C'', \geq, C' \rangle$);

refadd is the combination of addjoin and refine.

Figure 2 illustrates these operators. It shows two ontology fragments, left and right, with some assertions of disjointness (edges labelled \perp) and subsumption (unlabelled edges). Disjointness assertions simply materialises what can be deduced from class definitions by agents knowing them. During a game, the orange triangle instance is drawn and each agent knows to which class it belongs in its ontology (orange arrows). They identify the red correspondence (marked $=$) as erroneous: it leads to the conclusion that this instance belongs to a class which is disjoint from the identified class. delete removes this correspondence. replace replaces it with the subsumption part of the correspondence (\leq) that has not been proved incorrect. add (addjoin) will add a (\geq) correspondence from a (the common) subsumer of C (and D). The refine operator will add (\geq) correspondences to the subsumees of C' not subsuming D' .

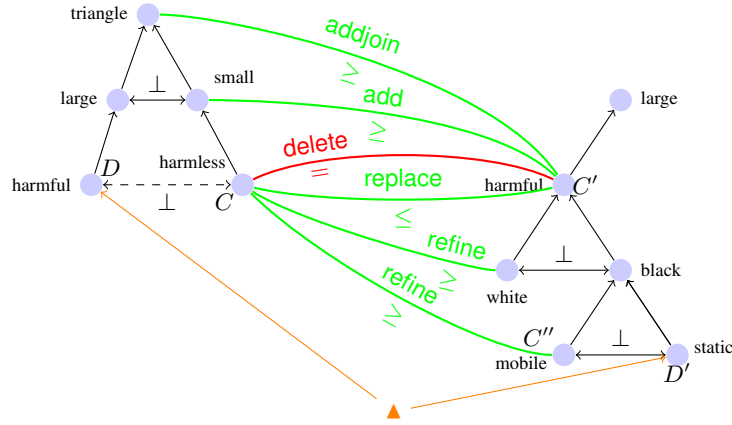


Fig. 2. Illustration on two ontology fragments of the effect of the various operators on the correspondence $\langle C, =, C' \rangle$ (adapted from [8]). \perp denotes class disjointness. The displayed object is characterised by the features: triangle, large, harmful, black and static.

3.4 Expansion and relaxation

We introduced two different modalities further refining agent behaviour [8]. In particular, the expansion modality permits agents to go beyond the initially provided alignments.

The **expansion** modality, when an operator removes a correspondence and is not able to replace it by another correspondence, adds a new random correspondence. This modality can be associated with any operator above.

Agents rely on random correspondences in the sense that there is nothing which guides them towards a particular correspondence. This modality is associated with a memory recording discarded correspondences so as to avoid regenerating them [8].

Moreover, correspondences that the agent can find entailed by its ontology and the current alignment or by a recorded correspondence will not be added.

The **relaxation** modality makes agents more curious: Agent A' has the opportunity to not use the most specific correspondence when answering to A , but to chose a more general one. This is alike someone answering a question with a general statement and not the most specific answer she knows (this is Latin music vs. this is Brazilian bossa nova vs. this is a Luiz Bonfá tune).

With relaxation, when encountering a correspondence that applies to the current instance, agents use this correspondence only with a specific probability, called immediate consumption probability, which also applies to the shadowed correspondences. As in [8], we use 80% immediate consumption probability (in 20% of the cases agents try to answer with less specific correspondences). Relaxation permits to detect incorrect correspondences that otherwise would never be detected because they are shadowed by more specific correct correspondences. This modality allows agents to reach 100% precision. In case there is no shadowed correspondence but the topmost one, agents use the applicable correspondence.

3.5 Measures

As agents play games, measures may be applied to the resulting situation (the current network of ontologies) and the evolution of the game. We use the same measures as in the initial experiments [7]:

- **Success rate** [16] characterised by the ratio of success over games played is the main measure.
- **Semantic precision and recall** [6] measures the degree of correctness and completeness of the resulting alignments with respect to the known correct reference alignments.¹
- **Incoherence rate** [12] measures the proportion of incoherent correspondences in alignments taken one by one.
- **Convergence** is the number of games taken to converge in all cases (it is an observed maximum, not an average) when the process converges.

The quality of an alignment is thus measured through its precision and recall, usually aggregated as F-measure. The incoherence rate is also a quality measure but it is negatively correlated with precision.

In [7], results were also compared with those of logical repair systems. Since, agents are allowed to generate new correspondences not entailed by the initial alignments, the comparison would not be meaningful any more. What would be relevant is to compare the agent results with that of ontology matching systems. However, this is more difficult to achieve than with logical repair systems which were relying on logic only. Indeed, ontology matching can take advantage of various features found in real-world ontologies that are not available in our synthetic ontologies.

¹ Contrary to classical precision and recall, it is not possible to deduce them from the numbers given hereafter.

In the following, we will test hypotheses by performing experiments on our *Lazy lavender* platform². We use only 4 agents to be able to compare with previous result and due to the complexity incurred in our setting by multiplying agents: this also increases ontology complexity. Experiments are run over 10000 random games. All results are the average of 10 runs (usually not the same ones).

The experiments are run with all six presented operators. We only deal with them under the expansion modality which is the only one generating correspondences. Hence, we compare with both expansion and expansion+relaxation modalities (see summary in Figure 7).

4 Starting empty

The expansion modality enables agents to introduce new correspondences. The question arises of whether this is sufficient to create the alignments in the first place. It is our first hypothesis (H1), that agents with expansion modality can reach the same quality if they start with an empty network of alignments.

4.1 Experimental setting: empty network

To test this, we experiment with agents starting with empty alignments. These alignments are not truly empty, they contain an equivalence correspondence between the topmost class of both ontologies. This materialises the idea that agents model the same domain. This is not a particularly strong assumption.

4.2 Results and discussion

The results are found in the third section (empty) of Table 1.

The exact same result is actually obtained by all adaptation operators, independently from the use of expansion and relaxation. This explains why the third section of Table 1 is simply reduced to one line. The resulting network is made of the initial empty alignments, which provides a coherent (incoherence=0.) and correct (precision=100%) alignment which is far from complete (recall=7%).

If agents start with empty alignments, then communication never fails—it can use the correspondence between the topmost classes which are correct—and the operators are never triggered, leading to no generation of new correspondences. So, agents are not able to evolve their knowledge.

This shows that Hypothesis H1 is not valid.

5 Generative modality: trying useful correspondences

To address this problem it is necessary to enhance agent behaviours. We achieve this by introducing a new modality, called generative, which allows agents, when they realise

² <http://lazylav.gforge.inria.fr>

		Network and Success		Incoherence	Semantic	Semantic	Semantic			
		operator	rate	Size	degree	Precision	F-measure	Recall	Convergence	
		reference	1.0	86	0.0	1.0	1.0	1.0	-	
		initial	0.24	54	0.34	0.11	0.20	0.89	-	
plain	expansion	delete	0.98	14	0.03	0.94	0.26	0.15	2003	
		replace	0.97	25	0.06	0.90	0.40	0.26	1664	
		refine	0.96	36	0.10	0.85	0.56	0.42	2798	
		add	0.96	59	0.23	0.69	0.67	0.65	4209	
		addjoin	0.98	51	0.20	0.74	0.65	0.57	2029	
		refadd	0.97	63	0.20	0.75	0.74	0.72	3203	
	exp+rel	delete	0.98	14	0.00	1.00	0.26	0.15	2851	
		replace	0.97	23	0.00	1.00	0.33	0.19	2916	
		refine	0.96	37	0.00	1.00	0.56	0.38	2610	
		add	0.95	49	0.00	1.00	0.64	0.47	7202	
		addjoin	0.97	42	0.00	1.00	0.58	0.41	3393	
		refadd	0.96	64	0.00	1.00	0.82	0.70	6002	
	empty	<i>any</i>	1.0	12	0.0	1.0	0.13	0.07		
	generative	expansion	delete	0.94	61	0.12	0.84	0.70	0.59	9099
			replace	0.94	63	0.12	0.84	0.70	0.59	8857
			refine	0.94	71	0.15	0.81	0.78	0.75	7063
			add	0.91	75	0.22	0.71	0.76	0.82	9344
			addjoin	0.95	75	0.21	0.75	0.78	0.81	6230
refadd			0.95	78	0.21	0.73	0.79	0.86	5840	
exp+rel		delete	0.94	52	0.00	1.00	0.60	0.43	8973	
		replace	0.94	57	0.00	1.00	0.65	0.48	6782	
		refine	0.93	63	0.00	1.00	0.77	0.62	6092	
		add	0.91	68	0.00	1.00	0.78	0.64	9907	
		addjoin	0.95	68	0.00	1.00	0.79	0.65	9501	
		refadd	0.94	74	0.00	1.00	0.87	0.76	8495	
generative+empty		expansion	delete	0.94	72	0.13	0.85	0.71	0.61	9199
			replace	0.94	71	0.12	0.84	0.70	0.59	7305
			refine	0.94	78	0.15	0.82	0.77	0.72	7458
			add	0.94	83	0.16	0.82	0.77	0.72	9173
			addjoin	0.96	80	0.16	0.81	0.78	0.74	5910
			refadd	0.95	82	0.18	0.80	0.80	0.80	7391
	exp+rel	delete	0.94	61	0.00	1.00	0.64	0.47	9504	
		replace	0.94	62	0.00	1.00	0.63	0.46	6997	
		refine	0.94	72	0.00	1.00	0.76	0.61	6798	
		add	0.93	76	0.00	1.00	0.78	0.63	8562	
		addjoin	0.95	75	0.00	1.00	0.78	0.64	9412	
		refadd	0.94	79	0.00	1.00	0.84	0.73	7078	

Table 1. Measures with the reference and initial network of ontologies, and those obtained by the 6 adaptation operators (delete, replace, refine, add, addjoin, refadd) with expansion, relaxation and generative modalities starting with random or empty networks of alignments [4 agents; 10 runs; 10000 games].

that they do not have a correspondence for a given object, to generate and immediately use a new correspondence that can apply to this object.

Because this changes the behaviour of the agents, we first run the modality under the same condition as the initial experiments (§3, [8]), i.e., starting with random alignments. This allows to be able to compare the impact of starting with empty correspondences (§6) on the same basis.

Our hypothesis is that (H2) the generative modality should not bring different results than the expansion modality which is already able to generate correspondences.

5.1 Generative modality

The generative modality applies when the only applicable correspondence for a given instance (\blacktriangle) is one of the topmost correspondence. Then the agent will generate a new correspondence that applies to this object. For that purpose, it will pick up at random a class (D) subsuming or equal to the class (C) of the instance and create a correspondences $\langle D, \leq, C' \rangle$ in which C' is taken randomly in the other ontology.

Figure 3 shows this situation highlighting the areas from which the correspondence is taken at random.




Fig. 3. Generative modality. Because no correspondences but the topmost one is applicable for \blacktriangle , a new one (in red, dashed) is drawn at random between the two gray areas of the ontologies.

This modality applies before any failure occurs. It injects correspondences which are immediately consumed. During the game, three situations may happen:

- the correspondence is correct: then communication will succeed and the correspondence will be preserved;
- the correspondence is incorrect and communication fails: then the correspondence is discarded immediately;
- the correspondence is incorrect and communication succeeds: then the correspondence is preserved and may be detected as incorrect later on (this is the case of Figure 3).

As in the expansion modality, agents record discarded correspondences so that they do not generate them again.

5.2 Experimental setting

We run the experiments comparing the results obtained by starting with random alignments and using or not using the new generative modality.

5.3 Results

The final results are found in the second (non generative) and fourth (generative) sections of Table 1 as well as the centre plots of Figure 5.

The resulting network features higher size when using the generative modality (see Figure 4). This is particularly relevant with the operators not allowed to add many correspondences (delete, replace and refine).




Fig. 4. Size of the resulting networks of ontologies without (left) and with (right) the generative modality. Expanding modality is in plain; expanding+relaxing in dashed [4 agents; 10 runs; 10000 games].

The generative modality is slower to converge: this is quite natural as it tries more correspondences.

Precision still reaches 100% when using relaxation. It however decreases (and inconsistency increases) when not using it. This is especially true for (delete and replace). This can be related to more correspondences generated which end up shadowed by other correspondences.

Recall is definitively higher for all operators causing higher F-measure for all conditions.

5.4 Analysis

Contrary to our expectation (H2), the generative modality allows agents to achieve better results than plain operators with expansion.

It is especially better in terms of recall. This is due to its capability to test correspondences even when the expansion modality would not have been applied: the expansion

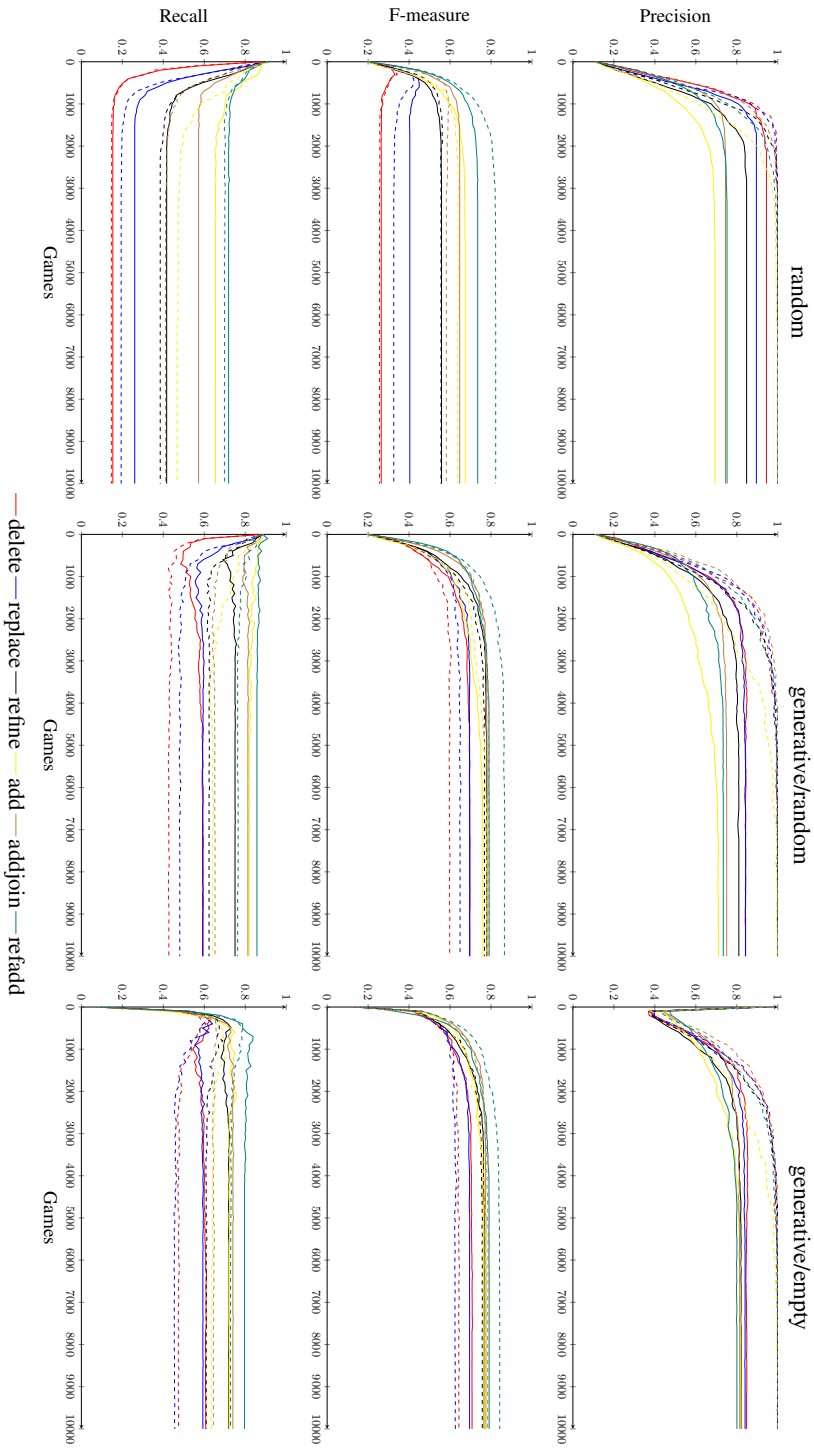


Fig. 5. Precision, F-measure and recall of the resulting networks of ontologies with random (left and centre) or empty (right) initial alignments and with (centre and right) or without (left) generative modality. Expanding modality is in plain; expanding+relaxing in dashed [4 agents; 10 runs; 10000 games].

modality adds correspondences, if possible, when the number of correspondences decreases, while the generative modality generates correspondences a priori and where they would be needed. This is especially visible for the worst operators (delete and replace), that had not many opportunities to remove correspondences.

The generative modality reaches a slightly lower precision and a far higher recall, which largely improves F-measure.

6 Generating from scratch

Results of Section 5 provide a new baseline: it is now possible to compare the behaviour of agents when they start with random or empty networks.

In this third experiment, the hypothesis (H3) is that agents starting from scratch are able to build alignments just like agents starting with random alignments.

6.1 Experimental setting

The experimental setting here is very simple as it combines the two previous ones: agents start with empty alignments and they use the generative modality. Experiments are run in the same conditions as before.

6.2 Results

The final results are found in the last section (generative+empty) of Table 1 and the right-hand side of Figure 5.

The size of the resulting networks follow the same curves after 500 games. Success rate are the same passed 200 games.

F-measure values seem even more concentrated than with random alignments.

6.3 Analysis

When using the generative modality, if agents start with empty alignments, they reach the same level of precision, recall and F-measure than if they would start with random alignments. We are talking of levels here because results are not exactly the same, but there is a lot of randomisation in these experiments. Looking at the raw figures of Table 1, there seems to be a slight increase of precision (except of course when relaxing is used) and a slight decrease in recall leading to an overall stable F-measure.

Hence, we do not think that this is strong enough to refute H3.

This means that the generation of initial random alignments is not necessary for agents to reach interoperability, they are able to generate tentative alignments and to discard them if they are wrong. Moreover, they will produce alignments with a commensurate quality, measured through F-measure, as if they were starting with a random network.

The centre and right curves of Figure 5 clearly differs as starting with empty alignments makes precision starting at 1. and steeply dropping, while recall starts near 0. and

sharply rises. On the contrary, starting with random alignments, precision is very low and rises steadily while recall is quite high and drops steeply before stabilising.

Figure 6 superimposes these curves to better assess their relations. It shows that after a few hundred games, the curves tend to have the same shape, though not exactly superimposed —again, keep in mind that the experiments are random.

Results could be further decomposed between the three first modalities (delete, replace, refine) and the three last ones (add, addjoin, refadd) which seems to benefit differently from the generative modality. For the three former ones, curves are very close indeed. This is different for the three latter ones: in case of relaxation, the curves are again nearly the same. However, without it, precision remains at a higher level and recall at a lower level when starting empty. This may be explained by the shadowing effect [8] which prevents from discarding incorrect correspondences that relaxation can debunk. Correspondences generated by the generative modality are never shadowed correspondences. Recall is higher because shadowed incorrect correspondences may entail correct correspondences accounted for by *semantic* precision and recall [6]. Both effects seems to compensate when computing F-measure which is globally the same.

Starting empty with the generative modality, agents achieve networks of alignments with more correspondences. It is not fully clear why this is the case.

Finally, although it seems that agents are actually traversing the whole search space, this is not the case. Indeed, they do not reach full recall and their networks of alignments are smaller than the reference network. This has been explained in [8] by the reverse shadowing effect which can clearly be observed in the case of relaxation: the size of alignments increases slightly, and they are 100% correct. However, their recall is lower. This counter-intuitive observation is related to enforcing functionality and preserving correct but not the most general correspondences.

7 Conclusion

Figure 7 synthesises all results. It presents generally valid evolution for relaxation, the generative modality and empty alignments (for refadd, but this also applies for the six operators). It show that the combination of relaxation, the generative modality and empty networks provides the best results.

The question of whether the expansion modalities introduced in [8] allows agents to create their own alignment from scratch can now be answered. The answer is that they cannot do this directly as they will not generate new correspondences. We thus introduced a game modality in which agents spontaneously test correspondences when no correspondence applies. This modality already improves the results when agents start with random alignments. With this modality, agents can start from empty alignments and will relatively quickly reach the same level of accuracy and coherence as with random alignments. This shows retrospectively that starting with random alignments is a realistic experimental condition, in the sense that it does not influence the stationary result in the end. This also shows that simulations can start from the origin.

This work is part of a wider investigation on experimental cultural knowledge evolution. Our first experiments aimed at developing local adaptation operators. We have




Fig. 6. Superimposition of the precision (left), F-measure (centre) and recall (right) curves for operators with expansion (red) and expansion+relaxation (blue) starting with random alignments (plain) and empty ones (dashed) [generative; 4 agents; 10 runs; 10000 games].




Fig. 7. Measures (above: Precision/F-measure/Recall; below: Size) obtained with the operator *refadd* in conjunction with the studied modalities (expansion, relaxation, empty start, generative). Red and green arrows present the predicted (non-)evolution of resulting networks. Yellow arrows present the observed evolution of measures when following an arrow of similar slope [*refadd*; 4 agents; 10 runs; 10000 games].

demonstrated that agents are able to acquire alignments from scratch through communication. We will now investigate how this can be achieved by a population of agents and to what extent they reach common knowledge that can be qualified of culture. We could then expose such populations of agents to changes in their environments and encountered population showing how adaptation turns into evolution.

Experiment material

Experiment records are available under the *Lazy lavender*² logbook at entries: [20170529-NOOR] [20170530-NOOR] [20170531-NOOR] [20170607-NOOR].

References

- [1] Karl ABERER, Philippe CUDRÉ-MAUROUX, and Manfred HAUSWIRTH, Start making sense: The Chatty Web approach for global semantic agreements, *Journal of Web Semantics* 1(1):89–114, 2003.
- [2] Michael ANSLOW and Michael ROVATSOS, Aligning experientially grounded ontologies using language games, in: *Proc. 4th international workshop on graph structure for knowledge representation, Buenos Aires (AR)*, pp. 15–31, 2015.
- [3] Manuel ATENCIA and Marco SCHORLEMMER, An interaction-based approach to semantic alignment, *Journal of Web Semantics* 13(1):131–147, 2012.

- [4] Paula CHOCRON and Marco SCHORLEMMER, Attuning ontology alignments to semantically heterogeneous multi-agent interactions, in: *Proc. 22nd European conference on artificial intelligence (ECAI), The Hague (NL)*, pp. 871–879, 2016.
- [5] Paula CHOCRON and Marco SCHORLEMMER, Vocabulary alignment in openly specified interactions, in: *Proc. 16th International conference on autonomous agents and multi-agent systems (AAMAS), Saõ Paolo (BR)*, pp. 1064–1072, 2017.
- [6] Jérôme EUZENAT, Semantic precision and recall for ontology alignment evaluation, in: *Proc. 20th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 348–353, 2007.
- [7] Jérôme EUZENAT, First experiments in cultural alignment repair (extended version), in: *Proc. ESWC 2014 satellite events revised selected papers, Lecture notes in computer science 8798*, pp. 115–130, 2014.
- [8] Jérôme EUZENAT, Interaction-based ontology alignment repair with expansion and relaxation, in: *Proc. 26th International Joint Conference on Artificial Intelligence (IJCAI), Melbourne (VIC AU)*, pp. 185–191, 2017.
- [9] 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, Ulm (DE)*, pp. 246–257, 2013.
- [10] Ernesto JIMÉNEZ-RUIZ, Terry PAYNE, Alessandro SOLIMANDO, and Valentina TAMMA, Limiting logical violations in ontology alignment through negotiation, in: *Proc. 15th Conferenve on Principles of Knowledge Representation and Reasoning (KR), Cape Town (ZA)*, pp. 217–226, 2016.
- [11] Christian MEILICKE, Alignment incoherence in ontology matching, PhD thesis, Universität Mannheim, 2011.
- [12] Christian MEILICKE and Heiner STUCKENSCHMIDT, Incoherence as a basis for measuring the quality of ontology mappings, in: *Proc. 3rd ISWC Ontology Matching workshop*, pp. 1–12, 2008.
- [13] Terry PAYNE and Valentina TAMMA, Negotiating over ontological correspondences with asymmetric and incomplete knowledge, in: *Proc. 14th International conference on autonomous agents and multi-agent systems (AAMAS)*, pp. 517–524, 2014.
- [14] Emanuel SANTOS, Daniel FARIA, Catia PESQUITA, and Francisco COUTO, Ontology alignment repair through modularization and confidence-based heuristics, *PLoS ONE* 10(12):1–19, 2015.
- [15] Luc STEELS, The origins of ontologies and communication conventions in multi-agent systems, *Autonomous agents and multi-agent systems* 1(2):169–194, 1998.
- [16] Luc STEELS, ed., *Experiments in cultural language evolution*, John Benjamins, Amsterdam (NL), 2012.
- [17] Cássia TROJAHN, Jérôme EUZENAT, Valentina TAMMA, and Terry PAYNE, Argumentation for reconciling agent ontologies, in: Atilla ELAI, Mamadou KONA, and Mehmet ORGUN, eds., *Semantic Agent Systems*, Springer, New-York (NY US), chap. 5, pp. 89–111, 2011.
- [18] Jurriaan VAN DIGGELEN, Robbert-Jan BEUN, Frank DIGNUM, Rogier VAN EIJK, and John-Jules MEYER, Ontology negotiation in heterogeneous multi-agent systems: The ANEMONE system, *Applied ontology* 2(3–4):267–303, 2007.