

ASMOV: Results for OAEI 2010

Yves R. Jean-Mary¹, E. Patrick Shironoshita¹, Mansur R. Kabuka^{1,2}

¹INFOTECH Soft, 1201 Brickell Ave., Suite 220, Miami, Florida, USA 33131

²University of Miami, Coral Gables, Florida, USA 33124
{reggie, kabuka}@infotechsoft.com

Abstract. The Automated Semantic Mapping of Ontologies with Validation (ASMOV) algorithm for ontology alignment has consistently been one of the top performing algorithms in the Ontology Alignment Evaluation Initiative (OAEI) contests. In this paper, we present a brief overview of the algorithm and its improvements, followed by an analysis of its results on the 2010 OAEI tests.

1 Presentation of the System

In recent years, ontology alignment has become popular in solving interoperability issues across heterogeneous systems in the semantic web. There exist many techniques to address this problem [1], differentiated by the way in which different ontology features are exploited. ASMOV, an algorithm that automates the ontology alignment process, uses a weighted average of measurements of similarity along four different features of ontologies, and obtains a pre-alignment based on these measurements. It then uses a unique process of semantic verification to ensure that the alignment does not contain semantic inconsistencies. In this manner, ASMOV was shown to produce the most coherent alignments of all systems tested in OAEI 2009 [3]. A more complete description of ASMOV is presented in [4].

1.1 State, Purpose, General Statement

ASMOV is an automatic ontology matching tool which has been designed in order to facilitate the integration of heterogeneous data sources modeled as ontologies. The current ASMOV implementation produces mappings between concepts, properties, and individuals, including mappings between object and datatype properties.

1.2 Specific Techniques Used

The ASMOV algorithm iteratively calculates the similarity between entities for a pair of ontologies by analyzing four features: lexical elements (id, label, and comments), relational structure (ancestor-descendant hierarchy), internal structure (property restrictions for concepts; types, domains, and ranges for properties; data values for individuals), and extension (instances of classes and property values). The measures

obtained by comparing these four features are combined into a single value using a weighted sum in a similar manner to [2]. These weights have been optimized based on the OAEI 2008 benchmark test results.

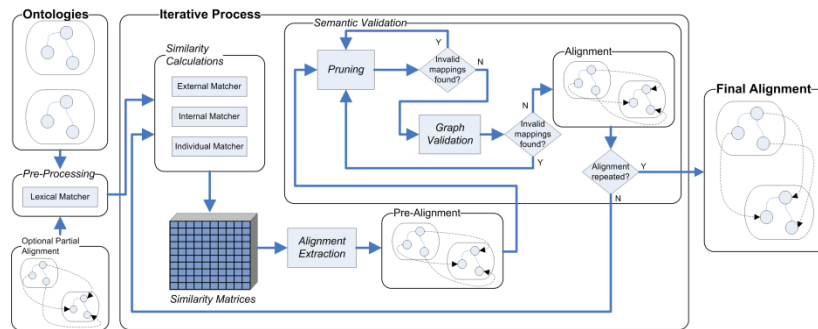


Fig. 1. The ASMOV Mapping Process

Fig. 1 illustrates the fully automated ASMOV mapping process, which has been implemented in Java. In the pre-processing phase, the ontologies are loaded into memory using the Jena ARP parser [5] and ASMOV's ontology modeling component. A thesaurus is optionally used to calculate the lexical similarities between each pair of concepts, properties and individuals. ASMOV can be configured to use either the UMLS Metathesaurus [6] or WordNet [7] in order to derive the similarity measures. If a thesaurus is not used, a text matching algorithm is used to compute the lexical distance. Following this, the similarities between pairs of entities along the relational structure, internal structure, and extensional dimensions are calculated, and overall similarity measures (or confidence values) are calculated for each pair. From these similarity measures, a pre-alignment is obtained by selecting the entity from one ontology with the highest similarity for a corresponding entity in the other ontology. A threshold of 0.1% is used to ignore spurious non-zero similarity measures.

This pre-alignment then goes through semantic verification, which detects semantically inconsistent mappings and their causes. These inconsistent mappings are removed from the pre-alignment and logged so that the algorithm does not attempt to map the same entities in a subsequent iteration; mappings are removed from the log of inconsistencies when the underlying cause disappears. Five specific types of inconsistencies are detected by ASMOV:

- Multiple entity correspondences, where the same entity on one ontology is mapped with multiple entities in the other ontology; unless these multiple entities are asserted to be equivalent, this type of mapping is unverified.
- Crisscross correspondences, where if a class c_1 in one ontology is mapped to some other class c_1' in the second ontology, a child of c_1 cannot be mapped to a parent of c_1' .
- Disjointness-subsumption contradiction, where if two classes c_1 and c_2 are disjoint in one ontology, they cannot be mapped to two other classes c_1' and c_2' in the second ontology where one is subsumed by the other. This also

applies to the special cases where c_1' and c_2' are asserted equivalent, or where they are identical.

- Subsumption incompleteness, if two classes c_1 and c_2 are mapped to two other classes c_1' and c_2' respectively in the second ontology, and if c_2 is subsumed by c_1 , then c_2' must be subsumed by c_1' , otherwise the correspondences are unverified. Similar incompleteness can be verified for the special case of equivalence.
- Domain and range incompleteness: if a class c_1 in one ontology is mapped to some class c_1' in the second ontology, and a property p_1 in the first ontology is mapped to some property p_1' in the second ontology, and if c_1 belongs to the domain (or range) of p_1 , then c_1' must belong to the domain (or, equivalently, range) of p_1' ,

Since OAEI 2009, ASMOV has been improved in three important respects, generally related to the new instance matching tests. The algorithm has generally been enhanced to allow it to process certain property constructs introduced in OWL 2, especially irreflexive and asymmetric properties. A procedure for disk-based storage of intermediate results has been implemented, allowing the algorithm to handle larger ontologies, although the ontology itself still needs to reside in memory. Further, we have improved the ability of ASMOV to use reasoning enabled by OWL in order to find semantically relevant matches. In particular, we have improved the verification of disjointness between domains and ranges of properties, and we also have included verification of functional properties.

1.3 Adaptations Made for the Evaluation

No special adaptations have been made to the ASMOV system in order to run the 2010 OAEI tests. The stop criterion for ASMOV was established as a multiple-alignment change threshold. For situations where both ontologies have more than 500 concepts, this threshold was set at 1% over three consecutive alignments; otherwise, it was set at 0% over two consecutive alignments. Although the rules of the contests stated that all alignments should be run from the same set of parameters, it was necessary to change two parameters for the anatomy tests. These parameters relate to the thesaurus being used (UMLS instead of WordNet) and to the flag indicating whether or not to use ids of entities in the lexical similarity calculations.

1.4 Link to the ASMOV System

The ASMOV system (including the parameters file) can be downloaded from <http://www.infotechsoft.com/products/asmov.aspx>.

1.5 Link to the Set of Alignments Produced by ASMOV

The results of the 2010 OAEI campaign for the ASMOV system can be found at <http://www.infotechsoft.com/products/asmov.aspx>.

2 Results

In this section, we present our comments on the results obtained from the participation of ASMOV in the five tracks of the 2010 Ontology Alignment Evaluation Initiative campaign. Tests were carried out on a PC running FreeBSD over VMware with two quad-core Intel Xeon processor (1.86 GHz), 8 GB of memory, and 2x4MB cache, with a Web service interface to run with the SEALS environment where required.

2.1 Benchmark

The OAEI 2010 benchmark tests have been divided by the organizing committee in eleven levels of difficulty; we have added one more level to include the set of 3xx tests, which have been included in the benchmark for compatibility with previous years. In Table 1, we present the results of these tests in comparison with those obtained in 2009 [8], where ASMOV was found to be one of the two best performing systems [3]. As can be seen, the precision, recall, and F1 measure for the entire suite of tests shows that ASMOV 2010 achieves 98% precision and 89% recall, and an F1 measure of 93%, which represents a 2% improvement over the 2009 version.

The accuracy of ASMOV in the benchmark tests is very high, especially for the lowest levels of difficulty. It is particularly noteworthy that improvements in both precision and recall were obtained especially at higher levels, with the largest improvement within level 10, the most difficult, and with significant improvements at levels 7 through 9 and at the 3xx tests. We believe that these improvements have come specifically through the enhancement of the procedures for utilizing domain and range information in the calculation of internal structure similarity, and through the correction of coding errors and deficiencies. In addition, some of this improvement can be attributable to improvements made in the gold standard.

Table 1. Benchmark test results for ASMOV version 2009 and version 2008

Level	ASMOV 2010			ASMOV 2009		
	Precision	Recall	F1	Precision	Recall	F1
0	1.00	1.00	1.00	1.00	1.00	1.00
1	0.99	1.00	0.99	1.00	1.00	1.00
2	1.00	0.99	0.99	1.00	0.99	0.99
3	0.99	0.98	0.98	0.99	0.98	0.98
4	1.00	0.98	0.99	0.99	0.98	0.98
5	0.99	0.94	0.96	0.97	0.93	0.95
6	0.98	0.90	0.94	0.95	0.89	0.92
7	0.98	0.87	0.92	0.93	0.83	0.88
8	0.98	0.77	0.86	0.90	0.71	0.79
9	0.97	0.64	0.77	0.83	0.48	0.61
10	0.90	0.29	0.44	0.40	0.04	0.07
3xx	0.88	0.84	0.86	0.81	0.82	0.81
All	0.98	0.89	0.93	0.95	0.87	0.91

2.2 Anatomy

For the anatomy track, ASMOV uses the UMLS Metathesaurus [6] instead of WordNet in order to more accurately compute the lexical distance between medical concepts. In addition, the lexical similarity calculation between concept names (ids) is ignored as instructed by the track organizers. ASMOV produces an alignment for all four subtasks of this track; the SEALS platform provides accuracy measurements for the first three subtasks.

1. *Optimal solution*: The optimal solution alignment is obtained by using the default parameter settings of ASMOV. The accuracy figures obtained from SEALS indicate precision of 79.9% and recall of 77.2%, resulting in overall F1 of 78.5%; these figures are a distinct improvement over the results obtained in 2009.
2. *Optimal precision*: The alignment with optimal precision is obtained by changing the threshold for valid mappings from 0.1% to 30%. The result is that precision increases to 86.5%, while recall decreases to 75.7%. F1 measure is 80.8%, which is higher than our optimal solution, indicating that the use of a higher threshold for ASMOV should be studied more closely.
3. *Optimal recall*: To improve recall, this time ASMOV made use of the annotation property `hasRelatedSynonym` included in the ontologies, to signify synonyms. It should be emphasized that this property is not included in the optimal solution because annotation properties do not have established semantics in OWL, and therefore it would not be possible for a computer to automatically understand that this property actually lists synonyms. The results from SEALS indicate that ASMOV found a total of 1521 alignments with precision of 71.7% and recall of 79.2%, resulting in F1 of 75.3%.
4. *Extended solution*: With a partial alignment given as input, the resulting alignment contained all mappings in the partial plus an additional 480 mappings.

2.3 Conference

This collection of tests dealing with conference organization contains 16 ontologies, of which at least one contains constructs specific to OWL 2. ASMOV is able to generate all 120 potential alignments from those ontologies.

Our analysis of the preliminary results obtained in running ASMOV against these ontologies showed a large number of erroneous matches due to incompleteness in our processing of disjointness between property domains and ranges. Specifically, our

Table 2. Results for Conference Test

F1	cmt	confer.	Confof	edas	ekaw	iasted	sigkdd
cmt		0.476	0.378	0.556	0.437	0.364	0.541
confer.			0.718	0.453	0.451	0.286	0.500
confof				0.549	0.681	0.378	0.357
edas					0.529	0.386	0.510
ekaw						0.348	0.368
iasted							0.481

previous versions of ASMOV only verified whether disjointness axioms existed in the asserted domain and range classes. We have now expanded ASMOV to verify any inferred disjointness between domains or ranges based on asserted disjointness within the subsumption hierarchy. Table 2 shows the results that were obtained by running the test through the SEALS platform.

2.4 Directory

Both the “small task” and the “single task” modalities were completed using ASMOV. The “small task” modality comprised 4639 tasks. We compared the results obtained this year against those obtained in 2009, where ASMOV was the best performing system [3]. We found a large degree of agreement, measured as 97% F1. We believe that the difference should result in improvement in the performance of ASMOV over 2009. The “single task” modality consisted of the alignment of a source ontology with 2854 classes, against a target ontology with 6555 classes. ASMOV found a total of 3347 mappings, with a large number of source ontology classes mapped to multiple target ontology classes.

2.5 Instance Matching

The application of ASMOV to the new set of IIMB instance matching tests results in precision of 86%, recall of 82%, and F1 measure of 84% for the small test, and precision of 85%, recall of 80%, and F1 measure of 82% for the large test.

The results of running the persons and restaurants (PR) tests in the SEALS platform are shown in Table 3. We noted the following issues:

- Some conflicts exist between URIs in the TBox (the description ontology) and the ABox. For example, the namespace URI for ontology_people1.owl in the person1 test was http://www.okkam.org/ontology_people1.owl in the TBox but http://www.okkam.org/ontology_person1.owl in the ABox. We manually corrected the TBox file where these differences were found to enable retrieval of the descriptions of the classes and properties used in the ABox.
- The gold standards for these tests only contained instances of the class “Person” in the person1 and person2 tests, and of the class “Restaurant” in the restaurant test. Running ASMOV in standard fashion produces alignments of instances of other classes such as “Address”; therefore, we restricted ASMOV to only find alignments of a pre-specified class in each test.
- The gold standard also contains mappings between instances that only match in one specific property, when other potential mappings contain matches in more properties. For example, in the Restaurant test, some matches in the gold standard are done exclusively over the “name” property, even if addresses and other property values are different.
- The gold standard also contains mappings between instances that have different values for functional properties. For example, the “surname” property is declared as functional for the class “Person” in the TBox, but two instances with “surname” property “carter” and “carcer” respectively are aligned in the gold

standard. The semantics of functional properties do not allow such an alignment, and ASMOV therefore rejects it. To test the effect of this, we ran ASMOV against the PR tests using and

Table 3. PR Instance Matching Results

	using TBox		ignoring TBox	
Person1	1.000	0.766	1.000	1.000
Person2	0.982	0.135	0.701	0.235
Restaurants	0.696	0.696	0.696	0.696

ignoring the TBox in the description ontology. As can be seen in Table 3, ASMOV obtains better results by ignoring the TBox in both Person tests.

3 General Comments

3.1 Comments on the Results

The current version of ASMOV has shown improvement overall in recall and F1 measure with respect to the results obtained last year in the benchmark tests. This is significant since the results in 2009 were already very high. The larger improvements have been obtained in the most difficult tests, showing the utility of the OAEI benchmarks in driving improvement of alignment algorithms. We have also been able to improve our accuracy in the benchmark, directory, and conference tasks. In the instance matching task we find some differences of interpretation with respect to the gold standard, specifically in terms of the semantics of certain properties.

3.2 Discussions on the Way to Improve ASMOV

ASMOV still needs to improve its ability to work with very large ontologies and resources. While some disk-based storage of partial results has been implemented, the entire contents of the ontologies still needs to be loaded in memory prior to performing the matching process. This needs to be further improved to use permanent storage in order to enable the alignment of very large ontologies. We also need to continue the implementation of the ability to infer assertions in order to utilize them for similarity measurement and semantic verification. In addition, we are also working in the improvement of the general scalability of the ASMOV algorithm for the processing of ontologies with a large number of entities. Finally, we need to reexamine the use of an appropriate threshold value to optimize accuracy.

3.3 Comments on the OAEI 2010 Test Cases

The new tests added to the OAEI 2010 contests provide important and welcome tools for the improvement of ontology matching systems. Most importantly, the instance matching task has been made significantly more challenging, allowing us to further refine and expand ASMOV to handle such alignments. Moreover, the availability of

an ontology in OWL 2 has allowed us to test some of the improvements made to ASMOV in light of the new standard. In addition, the ability to check accuracy using the SEALS system promises to help significantly in the debugging of our algorithms, once the technical problems with SEALS are resolved. Finally, the continuity in the benchmark, anatomy, and conference tracks allows us to evaluate the improvement of our algorithm and implementation as we proceed through its development.

One significant problem we found was the extended downtime encountered with the SEALS system. While it is understandable that some technical issues would be encountered, since this is its first deployment for OAEI, we found that SEALS hindered rather than helped in the process of debugging our algorithm and preparing our results. We trust and expect that many of these problems be resolved in the future, as SEALS promises to be a very useful tool for algorithm evaluation.

4 Conclusion

We have presented a brief description of an automated alignment tool named ASMOV, analyzed its performance at the 2010 Ontology Alignment Evaluation Initiative campaign, and compared it with its 2009 version. The test results show that ASMOV is effective in the ontology alignment realm, and because of its versatility, it performs well in multiple ontology domains such as bibliographic references (benchmark tests) and the biomedical domain (anatomy test). The tests results also showed that ASMOV is a practical tool for real-world applications that require on-the-fly alignments of ontologies.

Acknowledgments. This work is funded by the National Institutes of Health (NIH) under grant R43RR018667. The authors also wish to acknowledge the contribution of Dayana Paez, Michael Ryan, Ray Bradley, and Thomas Taylor of INFOTECH Soft, Inc.

References

1. Euzenat J and Shvaiko P. *Ontology Matching*. Springer-Verlag, Berlin Heidelberg, 2007.
2. Euzenat J. and Valtchev P. Similarity-based ontology alignment in OWL-lite. *In Proc. 15th ECAI*, Valencia (ES), 2004, 333-337.
3. Euzenat J, et.al. Results of the Ontology Alignment Evaluation Initiative 2009. http://www.dit.unitn.it/~p2p/OM-2009/oaei09_paper0.pdf. Accessed 1 Oct 2010.
4. Jean-Mary Y., Shironoshita E.P., Kabuka, M. Ontology Matching with Semantic Verification. *Journal of Web Semantics*. <http://dx.doi.org/10.1016/j.websem.2009.04.001>.
5. Jena from HP Labs <http://www.hpl.hp.com/semweb/>
6. Unified Medical Language System (UMLS) <http://umlsks.nlm.nih.gov/>
7. WordNet <http://wordnet.princeton.edu/>
8. Jean-Mary Y, Kabuka M. ASMOV: Results for OAEI 2009. http://www.dit.unitn.it/~p2p/OM-2009/oaei09_paper4.pdf.