Using AgreementMaker to Align Ontologies for OAEI 2009: Overview, Results, and Outlook*

Isabel F. Cruz, Flavio Palandri Antonelli, Cosmin Stroe,
Ulas C. Keles, and Angela Maduko

ADVIS Lab
Department of Computer Science
University of Illinois at Chicago
{ifc|flav|cstroe1|ukeles|maduko}@cs.uic.edu

Abstract. This paper describes our participation in the *Ontology Alignment Evaluation Initiative* (OAEI) 2009 with the AgreementMaker system for ontology matching, in which we obtained excellent results. In particular, we participated in the benchmarks, anatomy, and conference tracks. In the anatomy track, we competed against nine other systems in all four subtracks obtaining the best result in subtrack 3 and the second best result in subtracks 1 and 2. We were also first in finding the highest number of non-trivial correspondences. Furthermore, AgreementMaker came in first place among seven participants in the conference track and achieved the highest precision among all thirteeen participating systems in the benchmarks track. In addition to presenting this year's results, we give an overview of the AgreementMaker system, discuss ways in which we plan to further improve it in the future, and present suggestions for future editions of the OAEI competition.

1 Presentation of the system

As the Semantic Web evolves, more and more ontologies are being developed to describe conceptually several domains of interest. Ontology *matching* or *alignment*, which involves the task of finding correspondences called *mappings* between semantically related entities in two different ontologies, is needed to realize semantic interoperation and heterogenous data integration. A *matching* is a set of mappings established between two ontologies: the *source ontology* and the *target ontology*.

Automatic matching methods are highly desirable to allow for scalability both in the size and number of ontologies being aligned. Our collaboration with domain experts in the geospatial domain [7] has revealed that they value automatic matching methods, especially for ontologies with thousands of concepts. However, they want to be able to evaluate the matching process, thus requiring to be directly involved in the loop. Driven by these requirements, we have developed the AgreementMaker system[1] that integrates efficient automatic matching

* Research supported by NSF Awards ITR IIS-0326284, IIS-0513553, and IIS-0812258.

[1] www.AgreementMaker.org

strategies with a multi-purpose user interface and a module to evaluate matchings [3].

The problem of finding matchings is challenging on several counts. For example, a particular matching method may be effective for a given scenario, but not for others. Also, within the same scenario, the use of different parameters can change the outcome significantly. Therefore, our framework introduces a combined approach that takes advantage of several matching techniques focusing on different features of the ontologies and that allows for different parameters to be set. In particular, our architecture allows for serial and parallel composition where the output of one or more methods can be used as input to another method or several methods can be used on the same input and then combined. A set of mappings may therefore be the result of a sequence of steps called *layers*. The motivation behind this framework is to provide the capability of combining as many mapping layers as needed in order to capture a wide range of relationships between concepts in real-world scenarios [1]. There are parameters that can be defined for all methods, such as cardinality and threshold, whereas other parameters are method dependent. The parameter values can be set manually by the user or by automatic methods that take into account quality measures [2].

We have been developing AgreementMaker since 2001, with a focus on real-world applications [5, 8] and in particular on geospatial applications [4, 6, 7, 9–12, 16]. However, the current version of AgreementMaker and its implementation represents a whole new effort. Not only have we added significant new aspects to the system, but we also have almost completely reimplemented it in the last year. For example, in September of 2008 the previous implementation consisted of 9,000 lines of Java code, whereas in September of 2009 the new implementation had 29,000 lines.

The new AgreementMaker system [1–3] supports: (1) user requirements, as expressed by domain experts; (2) a wide range of input (ontology) and output (agreement file) formats; (3) a large choice of matching methods depending, on the different granularity of the set of components being matched (local vs. global), on different features considered in the comparison (conceptual vs. structural), on the amount of intervention that they require from users (manual vs. automatic), on usage (standalone vs. composed), and on the types of components to consider (schema only or schema and instances); (4) improved performance, that is, accuracy (precision, recall, F-measure) and efficiency (execution time) for the automatic methods; (5) an extensible architecture to incorporate new methods easily and to tune their performance; (6) the capability to evaluate, compare, and combine different strategies and matching results; (7) a comprehensive user interface that supports advanced visualization techniques and a control panel that drives all the matching methods and evaluation strategies; (8) a feedback loop that accepts suggestions and corrections by users and extrapolates new mappings.

1.1 State, purpose, general statement

AgreementMaker comprises a wide range of automatic matching algorithms called *matchers*, an extensible and modular architecture, a multi-purpose user interface, a set of evaluation strategies, and various manual (e.g., visual comparison) and semi-automatic features (e.g., user feedback loop). Given the automatic processing requirement imposed by OAEI, we could mainly make use of the first two features. In particular, we adopted seven different matchers for the competition and took advantage of the modular architecture to organize those matchers into four different matching layers. The evaluation techniques came into play only in the combination phase, to disambiguate the quality of the mappings to be selected.

Even though we could not take direct advantage of the user interface of AgreementMaker in the competition, we want to highlight its benefits prior to the competition. For example, the user interface can display any ontology (the largest ones we have tested have 30,000 concepts), therefore we were able to display the OAEI ontologies to investigate their characteristics (see Figure 1). In addition, we could test, tune, and evaluate both the individual matchers and the particular composition of matchers that we used in the competition.

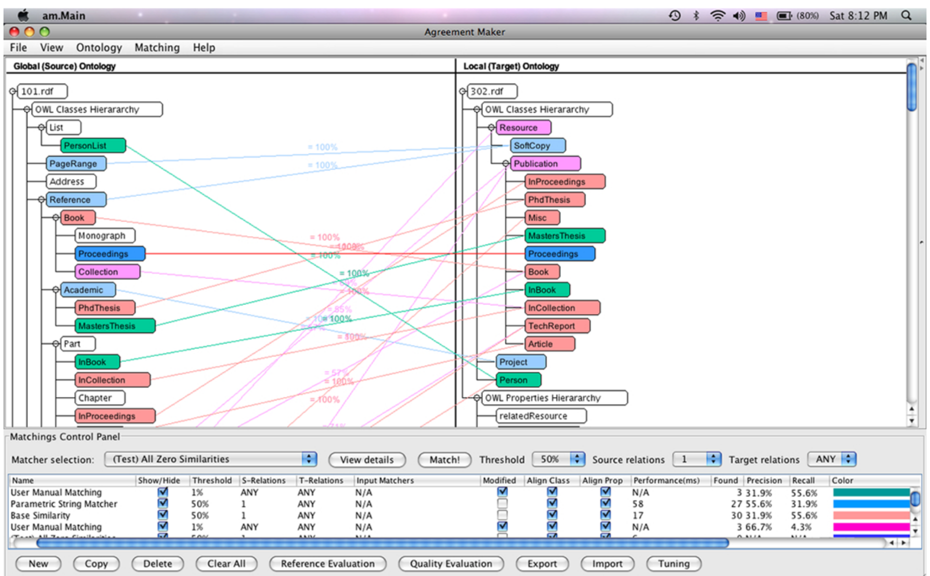

Fig. 1. Graphical User Interface of the AgreementMaker displaying ontologies from the benchmarks track.

1.2 Specific techniques used

For the OAEI 2009 competition, we have created a *stack of matchers*, shown in Figure 2, which are run on the input ontologies to compute the final alignment set.

First, three string-based techniques are independently run on the input ontologies: the *Base Similarity Matcher* (BSM) [7], the *Parametric String-based Matcher* (PSM) [2], and the *Vector-based Multi-word Matcher* (VMM) [2].

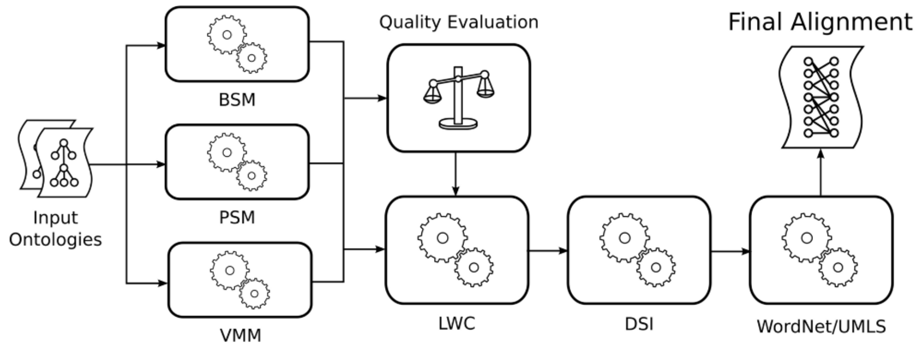

Fig. 2. AgreementMaker OAEI 2009 matcher stack.

BSM is a fundamental string-based matcher, which uses rule-based word stemming, stop word removal, and word normalization in order to find mappings. Going beyond the capabilities of BSM, PSM combines an edit distance measure and a substring measure in order to find mappings. Specifically for this campaign, PSM uses the following formula:

$$\sigma(a, b) = 0.6 * substring(a, b) + 0.4 * edit\_distance(a, b)$$

Our last string similarity matcher, VMM, compiles a *virtual document* for every concept of an ontology, then transforms the strings into TF-IDF vectors and computes the similarity using the *cosine similarity* measure.

After running the string matchers in parallel, their results are combined using the *Linear Weighted Combination* (LWC) matcher [2]. The LWC matcher uses the formula:

$$\sigma_{LWC}(a, b) = w_{BSM} * \sigma_{BSM}(a, b) + w_{PSM} * \sigma_{PSM}(a, b) + w_{VMM} * \sigma_{VMM}(a, b)$$

where the weights for each similarity are automatically calculated using the *local-confidence* quality measure. After the LWC matcher runs, we have a single, combined set of alignments that includes the best alignments from each of the string-based methods. The next matcher, the *Descendant's Similarity Inheritance* (DSI) [7] matcher, is a structure-based matcher that considers the ancestors of the concepts in a mapping in order to increase the similarity of the mapping. The DSI matcher is based on the following heuristic: if two nodes are matched with high similarity, then the similarity between the descendants of those nodes should increase. New mappings are created by the DSI matcher when the similarity of a mapping is increased beyond the threshold established for that matcher. The last step uses a lexical matcher, which considers not only the terms in an ontology, but also the synonyms of those terms as provided by a thesaurus (e.g., WordNet or UMLS).

In order to take advantage of the unique nature of the conference track, we performed an extra computation step, which we used in a new configuration of AgreementMaker called AgreementMakerExt. The OAEI 2009 matcher stack described above considers only two ontologies at a time. In order to expand this consideration, we have added a step that tries to take advantage of the transitivity between ontology mappings. We call this computation the *conflict resolution* step.

As shown in Figure 3, we consider two ontologies O_A and O_B, which have a mapping between them denoted $m_{A \leftrightarrow B}(a_i \in O_A, b_j \in O_B)$, given that concept $a_i \in O_A$ has been matched to concept $b_j \in O_B$. We then consider a third ontology O_C such that concept $a_i \in O_A$ is mapped to some concept $c_k \in O_C$ by mapping $m_{A \leftrightarrow C}(a_i \in O_A, c_k \in O_C)$. We also identify a mapping $m_{B \leftrightarrow C}(b_j \in O_B, c_h \in O_C)$ if there exists a concept $c_h \in O_C$ that matches $b_j \in O_B$. Note that $m_{A \leftrightarrow C}$ and $m_{B \leftrightarrow C}$ may point to different concepts in O_C (i.e., $k \neq h$).

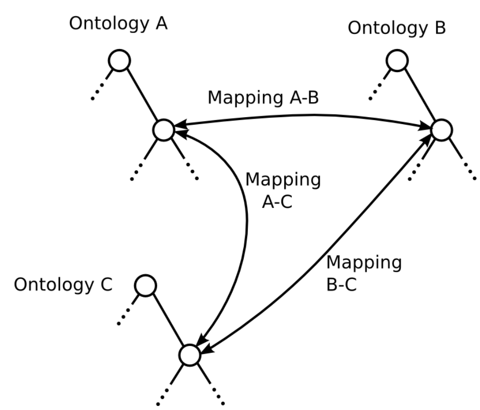

Fig. 3. Conflict resolution using a rating system.

We now implement a rating system. If $m_{A \leftrightarrow C}$ and $m_{B \leftrightarrow C}$ both map to the same concept in O_C (i.e., $k = h$), we increment the rating of all three mappings by 1. If $m_{A \leftrightarrow C}$ or $m_{B \leftrightarrow C}$ does not exist, we decrement the rating of any existing mappings by 1. Likewise, if $m_{A \leftrightarrow C}$ and $m_{B \leftrightarrow C}$ exist, but map to different concepts in O_C (i.e., $k \neq h$), we decrement the rating of all three mappings. This rating is performed for all the mappings between all the ontologies. Finally, we then sweep through the rated mappings and modify the alignments between any two ontologies to choose the mappings that have been rated the highest, resolving any conflicts by choosing the mappings with highest similarity.

1.3 Link to the set of provided alignments (in align format)

AgreementMaker alignment sets for OAEI can be found at
http://www.AgreementMaker.org/OAEI09_Results.zip.

2 Results

In this section, we present the results obtained by AgreementMaker in the OAEI 2009 competition. AgreementMaker participated in three tracks: benchmarks, anatomy and conference. Tests were carried out on a PC running Microsoft Windows Vista 64-bit with Intel Core 2 Duo 2.10 GHz processor and 4 GB RAM.

2.1 Benchmarks

In the benchmarks track, AgreementMaker uses the matchers described in Section 1.2. However, none of the lexical matchers was used in this track. The source ontology is compared to 111 ontologies (including the source ontology) in the same domain (bibliographic references). These ontologies can be grouped into three categories. We describe next the results obtained in each of these three categories as well as the overall results.

Concept test cases (1xx) There are four test cases in this category. AgreementMaker aligned the concept test cases with precision and recall equal to 98%.

Systematic test cases (2xx) For the systematic test cases, AgreementMaker achieved an average precision equal to 98% and an average recall equal to 60%. The average recall is lowered by the results of test cases in which the labels are scrambled. This is due to AgreementMaker's dependence on string mappings in order to find mappings based on structure. AgreementMaker achieved a precision in the range 94% to 100% and a recall in the range 85% to 100% in the systematic test cases in which labels are not scrambled.

Real ontology test cases (3xx) For the four real ontology test cases, AgreementMaker achieved an average precision equal to 92% and an average recall equal to 79%. Precision varied between 83% and 100% while recall varied between 60% and 95%.

Overall The overall results for all the categories place AgreementMaker first with precision equal to 99% and eighth with recall equal to 62% among thirteen participants.

2.2 Anatomy

The anatomy track of OAEI 2009 consists of finding alignments between two large real-world ontologies that are part of Open Biomedical Ontologies (OBO): the adult mouse anatomy (part of the Mouse Gene Expression Database) with 2744 classes and the human anatomy (part of the National Cancer Institute thesaurus) with 3304 classes.

This track consists of four subtracks. AgreementMaker entered all four subtracks using the UMLS Metathesaurus as background knowledge as well as the other modules in the OAEI 2009 matcher stack (see Figure 2).

The reference alignment contains 1523 mappings. Of these mappings only the 988 mappings that form part of the partial reference alignment for subtrack 4 are known. We note, however, that most of those mappings (934) are the "trivial" mappings that can be found by simple string comparison techniques. Therefore, the most important challenge is in finding the non-trivial mappings. The OAEI 2009 competition has released recall values for the non-trivial mappings for subtracks 1 and 3, and named this measure *recall+*. We describe next our results in the four subtracks.

Subtrack 1 In this subtrack, participants are asked to *maximize F-measure*. AgreementMaker used a threshold equal to 0.60 and obtained an F-measure equal to 83.1%, ranking second among ten participants and just short of the first ranked system (SOBOM) with F-measure equal to 85.5% and with a wider distance to the third ranked system (RiMOM) with F-measure equal to 79.3%. AgreementMaker obtained precision equal to 86.5% and recall equal to 79.8%, which was the highest recall value of all ten participants. AgreementMaker also ranked first in recall+, which was equal to 48.9%. The runtime was 23 minutes.

Subtrack 2 In this subtrack, participants are asked to *maximize precision*. AgreementMaker used a threshold equal to 0.75 and obtained precision equal to 96.7%, ranking second among seven participants and just short of the first ranked system (DSSim) with precision equal to 97.3% and with a wider distance to the third ranked system (TaxoMap) with precision equal to 95.3%. The runtime was 25 minutes.

Subtrack 3 In this subtrack, participants are asked to *maximize recall*. AgreementMaker used a threshold equal to 0.35. This choice of threshold combined with the UMLS module that was used to provide background knowledge resulted in our system having the highest recall among seven participants. AgreementMaker achieved recall equal to 81.5%. The second ranked system (Lily) had recall equal to 77.4%. AgreementMaker also ranked first in recall+, which was equal to 55.3%. The runtime was 17 minutes.

Subtrack 4 In this subtrack, participants are asked to *maximize precision, recall, and F-measure* using the mappings in a *partial reference alignment*, which is provided. AgreementMaker obtained the highest increase in precision among the four participants in this subtrack (+12.8%), significantly higher than that of the second ranked participant (ASMOV, with +3.4%). However, for recall and F-measure it was last (-18.1% and -6.3%, respectively).

2.3 Conference

In this track, participants are asked to find all correct correspondences (equivalence and/or subsumption correspondences) in a collection of 15 ontologies that describe a domain associated with the organization of conferences. Participants need to compute the set of mappings for each pair of ontologies. We note that for two ontologies O_i and O_j, $i \neq j$, once matching $M(O_i, O_j)$ is computed, the symmetric matching $M(O_j, O_i)$ need not be computed. Therefore, 105 alignment files need to be computed. In our case, the alignment files contained 2070 individual alignments in total.

We entered the competition with two different strategies. In one of them we used the OAEI 2009 matcher stack (see Figure 2), while in the other one, which we call AgreementMakerExt, we performed an extra computation step, which we described in Section 1.2. This step allows for more than two ontologies to be considered at a time by taking advantage of transitivity among ontology mappings. We call this computation the *conflict resolution* step.

From the "evaluation based on reference alignment" we see that AgreementMaker did very well overall. The alignments obtained by AgreementMaker with a threshold equal to 0.75 were the best among the seven participating systems, with precision equal to 69%, recall equal to 51%, and F-measure equal to 57%. The results obtained by AgreementMakerExt were also good, but the conflict resolution step reduced precision (to 54%), which led to a reduction of F-measure (to 51%). Since our system does not produce subsumption relations, it could not be evaluated on "restricted semantic precision and recall". Finally, in the "evaluation based on manual labeling", which rates how well the certainty of a system correlates with the correctness of the mappings, $80\% \pm 6\%$ of the mappings that were rated by AgreementMaker with a similarity equal to 1.0 were correct.

2.4 Comments on the obtained results

Benchmarks. Although AgreementMaker achieved the highest precision (99%) among all thirteen participating systems, it was less successful in terms of finding certain kinds of mappings, thus leading to less good recall (62%). This is because our structural techniques depend on lexical mappings that need to be found previously. When there were no lexical similarities, as was the case with some of the systematic test cases (2xx) where textual information was randomly modified, structural similarities were not found.

Anatomy. In subtracks 1-3, the most difficult task consists of finding those mappings that are non-trivial as observed in Section 2.2. Even so, AgreementMaker did very well in these subtracks having achieved first place in subtrack 3, second place in the remaining two, and first place in recall for non-trivial mappings. The key to further improvement relies on new techniques for finding other non-trivial correspondences.

In subtrack 4, as indicated in the work by Lambrix and Liu [14], partial reference alignments can be used at several points in the alignment process. We used the partial reference alignment that was provided in two ways:

1. To partition the ontologies into mappable parts so that every concept in the source ontology is not compared to every concept in the target ontology. We were able to reduce the running time of our algorithms by about 75%.
2. To remove mappings that are considered incorrect. Once the mappable parts are created, we assume that given a mappable part in the source ontology

and its corresponding mappable part in the target ontology, concepts at the same depth in the hierarchy match in the two mappable parts. We observe that this is especially true for ontologies that have similar structure.

Finally, partial reference alignments may be used to add undiscovered mappings to the final alignment results. This third aspect of using the partial reference alignment presents the most difficult challenges. We have not yet been able to implement a satisfactory method for accomplishing this third task. We hope to investigate this problem in future work.

Conference. AgreementMaker did very well on the conference track. AgreementMakerExt also did well in spite of the observed decrease in precision. In fact, the conflict resolution step decreased precision, while keeping recall almost the same. This leads us to infer that the conflict resolution step added wrong mappings and removed some correct mappings. We note, however, that the official results for this track were obtained using a partial reference matching that is one fifth the size of the full reference matching. The conflict resolution step works globally, that is, it may improve results overall, but not necessarily for just a "slice" of the problem; we therefore conjecture that this could provide the justification for the decrease in precision.

As for improving on the obtained results, our system ranked first with precision equal to 69% and recall equal to 51%, considering a threshold equal to 0.75. Precision can be further improved by understanding which mappings were erroneously included in the alignments, thus requiring an investigation of every single mapping in the alignment. Unfortunately, without the reference alignment we can only make an educated guess about which mappings are correct or incorrect. As far as improving recall, it seems that there is semantic information that we are not considering when aligning the ontologies. This may have to do with the unique nature of the conference track, in that it considers 15 ontologies mapped against one another instead of the traditional two.

2.5 Proposed new measures

In the anatomy subtracks 2 and 3, the participants are asked to compute an alignment that maximizes respectively precision (subtrack 2) and recall (subtrack 3). However, results that are based solely on the maximization of precision or recall may not be conclusive. For instance, a system could easily produce an alignment with 100% precision by computing an empty set of mappings, while an alignment containing all possible correspondences would have a 100% recall. Therefore, we suggest a different ranking system based on the use of a properly configured F-measure. To define our proposal, we first consider the definition of F-measure.

Given a set of mappings M and a reference matching R, the *F-measure of M with respect to R* is given by the following expression:

$$\textit{F-measure}(M,R) = \frac{\textit{precision}(M,R)\cdot \textit{recall}(M,R)}{(1-\alpha)\cdot \textit{precision}(M,R) + \alpha \cdot \textit{recall}(M,R)}$$

The higher the value of $\alpha \in [0,1]$, the more important is precision with respect to recall. Generally, it is set to 0.5 to get the harmonic mean of the two measures. In order to rank the matching results of the anatomy subtracks 2 and 3, we propose that they should be measured with respect to F-measure (not precision for subtrack 2 and recall for subtrack 3). Therefore, α should be greater than 0.5 for subtrack 2 and lower than 0.5 for subtrack 3. The value for α could be chosen by considering a ranking among the results obtained for the anatomy subtracks 2 and 3 from previous years. Once that ranking is established, then the corresponding value of α would be given to the OAEI participants in future editions of the competition so that they can tune their methods for that particular value.

Finally, we want to point to the fact that AgreementMaker can be used to import the OAEI alignments computed by any matching system in order to evaluate precision, recall, and F-measure thus allowing for their direct comparison. In addition, AgreementMaker can evaluate structural discrepancy measures [13] and the *local-confidence* quality measure that we defined [2]. We further plan to implement the incoherence-based quality measure [15].

2.6 Discussions on ways to improve the proposed system

There are several directions that we would like to explore to improve AgreementMaker. For example, we want to add matchers that rely solely on the structure of the ontologies to find matchings. The DSI (Descendant's Similarity Inheritance) and SSC (Sibling's Similarity Contribution) structure-based matchers exploit the structure of the ontology by respectively considering the concepts that have as descendants or siblings the concepts being matched [7, 16]. However, they first rely on similarity values computed by string-based matchers. We hope to devise "pure" structure-based matchers that would work in the benchmarks track cases where AgreementMaker did not produce any mappings, even though the ontologies being matched are very similar structurally.

AgreementMaker was not able to fully exploit the unique nature of the conference track. One way to further improve our results in the conference track is to incorporate the capabilty of extending alignments over multiple ontologies, instead of considering only two ontologies at a time.

Finally, we will further explore how to use the provided partial reference alignment in subtrack 4 of the anatomy track. In particular, we encountered some false negatives due to the dissimilarity in structure of the two anatomy track ontologies. We hope to devise other techniques that circumvent this. In addition, we would like to use the partial reference alignment to discover nontrivial mappings.

However, we are not only focusing on automatic ontology matching. We believe that involving the user in the matching process is crucial in finding the mappings that are not found by automatic methods. By taking advantage of the multi-purpose user interface of the AgreementMaker, we have been working on a semi-automatic matching approach that ranks concepts according to their relevance and presents to users the top-k most relevant concepts together with the most likely mappings associated with them. In addition, our solution encompasses a feedback loop that extrapolates new correspondences and corrects wrong mappings.

3 Conclusions

In this paper we gave an overview of the new AgreementMaker system, which was developed in the last year, presented the results that this system obtained in the OAEI 2009 competition for the benchmarks, anatomy, and conference tracks, and discussed those results. We also proposed new measures for future OAEI competitions.

In the benchmarks track, AgreementMaker found alignments (a total of 111) for all cases. All those alignments were computed in less than 3 minutes with an overall precision equal to 99% (highest among 13 competing systems) and an overall recall equal to 62% (eighth place).

AgreementMaker participated in all four subtracks of the anatomy track, placing second in subtracks 1 and 2 and first in subtrack 3 among ten, seven, and seven participants, respectively. AgreementMaker also found the highest number of non-trivial correspondences. In the last subtrack, subtrack 4, it achieved the highest improvement in precision among four participants together with an improved execution time.

AgreementMaker was also very successful in the conference track: it achieved the best results among seven participants, with precision equal to 69% and recall equal to 51% for a threshold equal to 0.75.

Overall, AgreementMaker exhibited an excellent performance in the OAEI 2009 competition. However, the competition only tests the component of AgreementMaker that performs automatic matchings. The automatic matching capabilities of AgreementMaker are just a small part of a full framework for ontology matching, which also supports the visualization of ontologies and the evaluation of their matchings. Those matchings can also be produced manually, semi-automatically, or using an extrapolating mechanism that accepts input from users. Several of these components of AgreementMaker, even if not directly tested in the competition, were quite useful for "understanding" the ontologies and for the tuning and evaluation of the matching strategies. However, we believe that there is still room for improvement and we plan to continue our quest for efficiency and effectiveness in the ontology matching process.

References

1. Isabel F. Cruz, Flavio Palandri Antonelli, and Cosmin Stroe. AgreementMaker: Efficient Matching for Large Real-World Schemas and Ontologies. *PVLDB*, 2(2):1586–1589, 2009.
2. Isabel F. Cruz, Flavio Palandri Antonelli, and Cosmin Stroe. Efficient Selection of Mappings and Automatic Quality-driven Combination of Matching Methods. In *ISWC International Workshop on Ontology Matching*. CEUR-WS, 2009.

3. Isabel F. Cruz, Flavio Palandri Antonelli, and Cosmin Stroe. Integrated Ontology Matching and Evaluation. In *International Semantic Web Conference (Posters & Demos)*, 2009.
4. Isabel F. Cruz and Afsheen Rajendran. Exploring a New Approach to the Alignment of Ontologies. In *ISWC Workshop on Semantic Web Technologies for Searching and Retrieving Scientific Data*, volume 83 of *CEUR-WS*, pages 7–12, 2003.
5. Isabel F. Cruz and Afsheen Rajendran. Semantic Data Integration in Hierarchical Domains. *IEEE Intelligent Systems*, March-April:66–73, 2003.
6. Isabel F. Cruz, Afsheen Rajendran, William Sunna, and Nancy Wiegand. Handling Semantic Heterogeneities Using Declarative Agreements. In *ACM International Symposium on Advances in Geographic Information Systems (ACM GIS)*, pages 168–174, 2002.
7. Isabel F. Cruz and William Sunna. Structural Alignment Methods with Applications to Geospatial Ontologies. *Transactions in GIS, special issue on Semantic Similarity Measurement and Geospatial Applications*, 12(6):683–711, 2008.
8. Isabel F. Cruz, William Sunna, and Anjli Chaudhry. Ontology Alignment for Real-World Applications. In *National Conference on Digital Government Research (dg.o)*, pages 393–394, 2004.
9. Isabel F. Cruz, William Sunna, and Anjli Chaudhry. Semi-Automatic Ontology Alignment for Geospatial Data Integration. In *International Conference on Geographic Information Science (GIScience)*, volume 3234 of *Lecture Notes in Computer Science*, pages 51–66. Springer, 2004.
10. Isabel F. Cruz, William Sunna, Nalin Makar, and Sujan Bathala. A Visual Tool for Ontology Alignment to Enable Geospatial Interoperability. *Journal of Visual Languages and Computing*, 18(3):230–254, 2007.
11. Isabel F. Cruz, William G. Sunna, and Kalyan Ayloo. Concept Level Matching of Geospatial Ontologies. In *GIS Planet International Conference and Exhibition on Geographic Information*, 2005.
12. Isabel F. Cruz and Huiyong Xiao. Data Integration for Querying Geospatial Sources. In John Sample, Kevin Shaw, Shengru Tu, and Mahdi Abdelguerfi, editors, *Geospatial Services and Applications for the Internet*, pages 113–137. Springer, 2008.
13. Cliff Joslyn, Patrick Paulson, and Amanda White. Measuring the Structural Preservation of Semantic Hierarchy Alignment. In *ISWC International Workshop on Ontology Matching*. CEUR-WS, 2009.
14. Patrick Lambrix and Qiang Liu. Using Partial Reference Alignments to Align Ontologies. In *European Semantic Web Conference*, volume 5554 of *Lecture Notes in Computer Science*, pages 188–202, 2009.
15. Christian Meilicke and Heiner Stuckenschmidt. Incoherence as a Basis for Measuring the Quality of Ontology Mappings. In *ISWC International Workshop on Ontology Matching*, volume 431 of *CEUR-WS*, 2008.
16. William Sunna and Isabel F. Cruz. Structure-Based Methods to Enhance Geospatial Ontology Alignment. In *International Conference on GeoSpatial Semantics (GeoS)*, volume 4853 of *Lecture Notes in Computer Science*, pages 82–97. Springer, 2007.