MAMBA - Results for the OAEI 2015

Christian Meilicke

Research Group Data and Web Science
University of Mannheim, 68163 Mannheim, Germany
christian@informatik.uni-mannheim.de

1 Presentation of the system

Most matching systems implement their functionality as a sequential process. Such systems start with analyzing different types of evidence, in most cases with a focus on the involved labels, and generate, as an intermediate result, a set of weighted matching hypotheses. From the intermediate result a subset of the generated hypotheses is chosen as final output. The approach implemented in MAMBA differs significantly from this approach.

MAMBA[1] treats labels (and their parts) as well as logical entities (classes and properties) as first class citizens in an optimization problem. During the matching process MAMBA generates hypotheses about equivalences between labels and tokens, while at the same time mappings between concepts and properties are considered to be true and wrong. MAMBA uses Markov Logic [6] to define constraints that ensure that the underlying assumptions about equivalent tokens are always consistent and that dependencies between labels and entities described by these labels are taken into account. The approach implemented in MAMBA has been described in details in a paper [4] that can also be found in the proceedings of the Ontology Matching Workshop. To avoid redundancy, we omit a description of the underlying approach in this paper. Instead of that we comment on some results and discuss open issues.

MAMBA is available at http://web.informatik.uni-mannheim.de/mamba/. Note that MAMBA was developed with the motivation to illustrate the benefit of the approach roughly sketched in [3] and finally presented in [4]. Thus, MAMBA is not a general-purpose ontology matching system but a research prototype.

2 Results

2.1 Conference Track

The OAEI conference track was used as one of the main test sets used during the development and testing of MAMBA. The achieved results are shown in Table 1.

Comparing these results against the results of previous OAEI editions, MAMBA is always among the best two systems with respect to F-measure. Only the system YAM++ [1] achieved an F-measure of .71 (ra-2) and .74 (ra-1), which is a bit better than the results of MAMBA.

[1] MAMBA stands for **Ma**nnheim **M**atcher based on a **B**ilayered **A**pproach

Gold Standard	Precision	F-measure	Recall
ra-1	.80	.68	.59
ra-2	.83	.72	.63

Table 1. Results for the Conference track

2.2 Results for the other tracks

Due to the fact that MAMBA is currently only a research prototype mainly developed for testing the approach that we described in [4], we have not conducted many experiments on other data sets. However, we already know that MAMBA will probably not be able to match ontologies with more than 1000 concepts due to the underlying optimization problem. Furthermore, we made only a very quick test with the bibliographic benchmark, to ensure that the basic functionality of a matching system is implemented.

3 General comments

3.1 Comments on the Results

The results for the Conference track illustrate the benefits of the proposed approach. Note that we applied a very restrictive approach for computing the input similarities which are used as evidence for the equivalence hypotheses between the tokens. We used more or less the maximum of Levensthein similarity and Wu Palmer WordNet similarity together with a very simple method for generating similarities between pairs of tokens that contain abbreviations (e.g., *ProgramCommitteeMember* vs. *PCMember*). Most approaches use a richer set of method with a fine tuned aggregation method. Thus, we believe that the results of MAMBA can be improved by using better similarity measures.

We did not compute results for any other track. While we were mainly interested in understanding the impact of our new approach, we could spend only a limited time in checking whether MAMBA is capable of generating alignments for all kind of input ontologies that might differ in format and in the way how labels are used to describe the logical entities. Preliminary experiments with one test set from the benchmark series showed that MAMBA generates for these synthetic data sets only mediocre results.

The most critical issues are related to the runtimes of MAMBA. MAMBA will not terminate for ontologies with more than 1000 concepts. The optimization problem that needs to be solved is NP-hard. Note also the the runtime performance of MAMBA is even worse than the runtime performance of CODI [2], which also defines internally an optimization problem. Due to the two layers of tokens and entities, MAMBA translates a matching problem into a more complex problem with more variables and more constraints.

3.2 Improving the Approach

An additional amount of engineering work is required to make MAMBA more robust. There is a high chance that the current version contains several bugs that need to be detected via extensive testing. We know, for example, that complex domain and range restrictions are currently not correctly interpreted by MAMBA.

The runtime problems of MAMBA cannot be solved easily. We are currently using a stack of systems (Rockit [5], GUROBI), where each system is known to be one of the most efficient systems for solving the type of problems that MAMBA generates. Moreover, we apply already a specific technique to speed up the matching task, by first solving a relaxed version of the matching problem, which allows to solve the harder problem more efficiently.[2]

Our main motivation while developing MAMBA was to show the need for generating alignments that are consistent with respect to the corresponding assumptions about the meanings of the involved tokens. This general idea is not necessarily bound to the use of optimization techniques. Greedy techniques can also be used to ensure this special kind of label/entity alignment consistency. Indeed, such approaches have to be used to make the general idea applicable to matching larger ontologies as we find them in the Anatomy track or in the Large Biomedical track.

3.3 Comments on OAEI test cases

The availability of the OAEI test cases has revealed that MAMBA needs to be significantly improved to become a robust matching systems instead of being just a set of scripts that have been used to illustrate the benefits of a specific approach. We must admit that we underestimated the engineering work that is required to implement these improvements.

However, our sole focus on the conference track was mainly motivated by the fact that the conference track is the only track that has a manually generated, high quality gold standard that is at the same time easily understandable, while the ontologies are relatively expressive and differ partially in their modeling style. This real world scenario results in a great deal of non trivial mappings that our approach is designed to detect. For that reasons it would be a significant improvement if the OAEI would offer a second track that has a similar characteristic as the conference track.

4 Conclusion

MAMBA is our attempt to implement the approach described in [4] as a matching system. While we were able to generate good results for the test cases of the Conference track, we have not yet systematically tested the performance of MAMBA for the other tracks. We already know that MAMBA will not terminate in acceptable time for test cases with more than 1000 classes. Nevertheless, the good results that we achieved for the conference track might be a motivation to modify existing matching systems in a way that the resulting mappings are consistent with respect to the implicit assumptions regarding the equivalence of the involved tokens.

[2] Unfortunately, this approach is not even explained in [4]. Contact the author if you are interested in the details.

References

1. Ngo Duyhoa and Zohra Bellahsene. Yam++ results for oaei 2013. In *Proceedings of the 8th International Workshop on Ontology Matching (OM 2013)*, 2013.
2. Jakob Huber, Timo Sztyler, Jan Noessner, and Christian Meilicke. Codi: Combinatorial optimization for data integration–results for oaei 2011. *Ontology Matching*, 134, 2011.
3. Christian Meilicke, Jan Noessner, and Heiner Stuckenschmidt. Towards joint inference for complex ontology matching. In *AAAI (Late-Breaking Developments)*, 2013.
4. Christian Meilicke and Heiner Stuckenschmidt. A new paradigm for alignment extraction. In *Proceedings of the Tenth International Workshop on Ontology Matching (OM 2015)*, 2015.
5. Jan Noessner, Mathias Niepert, and Heiner Stuckenschmidt. RockIt: Exploiting parallelism and symmetry for map inference in statistical relational models. 2013.
6. Matthew Richardson and Pedro Domingos. Markov logic networks. *Machine learning*, 62(1-2):107–136, 2006.