

Experimenting with ontology distances in semantic social networks: methodological remarks

Jérôme David
INRIA & LIG, France
Montbonnot, France
Jerome.David@inria.fr

Jérôme Euzenat
INRIA & LIG
Montbonnot, France
Jerome.Euzenat@inria.fr

Jason J. Jung
Yeungnam University
Gyeongsan, Korea
jj2jung@gmail.com

Abstract—Semantic social networks are social networks using ontologies for characterising resources shared within the network. It has been postulated that, in such networks, it is possible to discover social affinities between network members through measuring the similarity between the ontologies or part of ontologies they use. Using similar ontologies should reflect the cognitive disposition of the subjects. The main concern of this paper is the methodological aspect of experimenting in order to validate or invalidate such an hypothesis. Indeed, given the current lack of broad semantic social networks, it is difficult to rely on available data and experiments have to be designed from scratch. For that purpose, we first consider experimental settings that could be used and raise practical and methodological issues faced with analysing their results. We then describe a full experiments carried out according to some identified modalities and report the obtained results. The results obtained seem to invalidate the proposed hypothesis. We discuss why this may be so.

Index Terms—Semantic social networks; ontology distance; ontology similarity; personal ontologies; experimental methodology.

I. INTRODUCTION

Semantic social networks are characterised by a social network, i.e., relations between people, involving ontologies [1], [2]. These ontologies may be used for categorising resources, such as pictures or music, or for describing relationships themselves. Using ontologies provide more expressiveness to social network users for communicating, e.g., expressing queries, as well as more flexibility since users may choose the ontology they want or design their owns. Semantic social networks may be analysed in three layers (Fig. 1) which are (i) a concept layer, (ii) a ontology layer, and (iii) a social layer.

We are not considering ontologies here as a monolithic reference classification agreed by a group. Our focus is rather on personal ontologies that people may build for personal purposes and that may be used for communicating between them. The goal of these people is not to build the ultimate ontology but rather to design an ontology that meet their needs and which may evolve if these needs evolve. These are true ontology because people make efforts to model their domain and sometimes to reuse existing modeling, even if these people do not feel expert in one domain nor in modeling ontologies.

A common hypothesis is that, in semantic social networks, the closer the ontologies used by people, the more likely they are to develop fruitful relations. This hypothesis may be

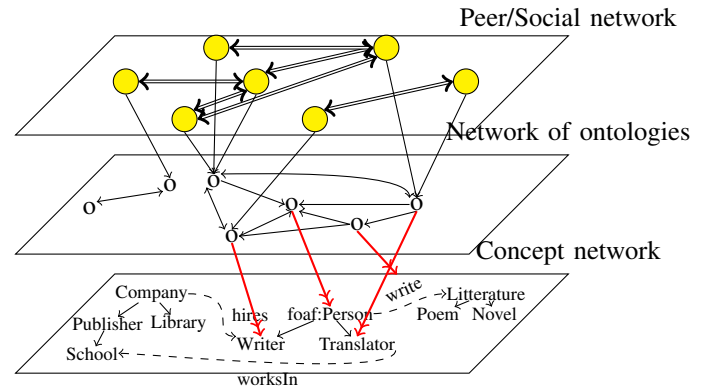


Fig. 1: A semantic social network (Inspired from [2]).

exploited to efficiently support social collaboration between people by estimating potential relationships between them and organizing communities over the social networks. Contrary to other approaches relying on relationships between users, such methods would take advantage of distances between user ontologies. Instead of clustering people who know each others, it would group people using similar ontologies assuming that they have some chance to share interests. The methodology is the following:

- 1) compute similarity between user ontologies,
- 2) cluster users from this similarity, e.g., by using hierarchical clustering algorithms.

Ontology distances and similarities in semantic social networks can be computed from:

- the ontology space (third layer),
- the alignment space (second layer), and
- a combination between both ontology space and alignment space.

In this study, we want to test this hypothesis and to evaluate measures for comparing ontologies in this context. We would like to evaluate the accuracy of various such measures in predicting affinity between users.

For that purpose, after considering the state of the art in semantic social networks (Sect. II) we consider the design space of experiments for evaluating the hypothesis (Sect. III) in order to illustrate the variability and difficulty of such experiments. Then we relate the implementation of such an

experiment (Sect. IV) and analyse its results (Sect. V).

II. RELATED WORK

Since [3] has shown how to build semantics-based social networks, there have been various applications, e.g., BibSter [4] and GrouPeer. In particular, [5], [6] proposes an elaborated method for identifying what types of semantics can be applied to relationships between people in Polyphonet. Most of these methods have focused on discovering co-occurrence patterns between users, and they have assumed that such co-occurrence patterns should be an important evidence implying social connections. While they focus on discovering common human behaviors, e.g., tagging and bookmarking, semantic overlay networks [7], [8] utilise auxiliary information to compare members. Ontologies themselves have not been exploited for finding relations between people. [4] went further by using similarity between ontology entities for routing queries to supposedly adequate peers, but did not use this for inferring affinities independently of queries.

Ontology-based overlay structures [1], [2], have been proposed to target user-created ontologies as well as semantic matchings between the ontologies. Hence, the semantic matchings between ontologies are assumed to increase the strength of social affinity between the corresponding people. However, this hypothesis has not been evaluated experimentally.

III. EXPERIMENT DESIGN

The goal of the experiment is to test the main hypothesis and to evaluate the accuracy of similarity measures. For that purpose, it is necessary to collect information about semantic social networks and affinity between people. This kind of information is not particularly easy to obtain: there are no widely used semantic social network yet, and it is difficult to obtain non biased affinity information.

A. Acquiring semantic social networks

The main problem we are facing is the acquisition of semantic social network data, i.e., linked (even loosely) sets of ontologies.

There are five possible strategies we can think of for acquiring the semantic social network.

1) *Strategy 1: Finding some existing semantic social network:* This is the most natural strategy since it would rely on real world data and not data made on purpose for the experiment.

- pros: This does not cost effort. The data is realistic.
- cons: No corpus readily available. There are no currently developed semantic social network.

2) *Strategy 2: Using collaborative ontology building systems:* Collaborative ontology building systems, e.g., Cupboard [9] supports online cooperative ontology management. They are aimed at ontology engineers and allow them for storing online and annotating the ontology they use. Using such systems would allow to have direct relation between people and real ontologies in a detailed way. On the other side, the relational aspect of the network would not be really developed.

- pros: It is maintained. All the ontologies of the web are there, through Watson.
- cons: Cupboard is for ontology engineers, not for social networks, so this may lead to non relevant results.

3) *Strategy 3: Developing a semantic social network:* We can invite some volunteers, e.g., students and team members, to use a semantic social network software (PicSter, BibSter) to perform an experiment. [2] already run small experiments under a dedicated peer-to-peer software (PicSter).

- pros: The experiment would be extensive.
- cons: The software on which these experiments were based is not anymore maintained. This require additional involvement from the subjects.

4) *Strategy 4: Using some existing social network data using tags instead of ontologies:* It may be possible to use classical social network data and to consider how this approximates semantic social networks. There is data available in which ontologies are replaced by rich set of tags (“folksonomies”). These sets of tags could be considered as ontologies without structures. This would still allow to use some measures on these sets of tags. We could also have trivial alignments on tags (based on WordNet for instance).

- pros: It will be easier to find data. The approach could prove useful in a concrete setting.
- cons: These are not real ontologies, so no semantic social networks.

5) *Strategy 5: Simulating social network data within a controlled experiment:* This is the usual procedure in experimental psychology: asking a set of people (subjects) to perform a task that is assumed to have the same characteristics as the tested task.

- pros: Relatively easy to perform;
- cons: The experiment is remote from real semantic social network use.

B. Acquiring affinity information

The second problem to address is to determine the relatedness between users. There are basically two ways:

- direct: Asking people to evaluate their affinity with other people.
- indirect: Inferring it from some cues. For instance, their tendency to annotate the same pictures.

We consider again four different strategy to acquire these data. We consider them under the constraint that we will ask people to annotate pictures with ontologies and that the choice of pictures may be used for assessing affinity.

1) *Strategy 1: Spontaneous affinity grouping:* This is the simplest strategy: Participants are asked to express their affinity, for instance on a scale. The same $|P|$ pictures are given to all the subjects.

- pros: All students have the same pictures.
- cons: It is very difficult to evaluate the validity of this affinity assessment.




Fig. 2: Distribution of pictures to subjects. Each leaf is a subject, each letter is a different picture.

2) *Strategy 2: Thematic grouping*: Participants are given different sets of pictures corresponding to different topics:

- pros: Topics will induce affinity.
- cons: We will observe the correlation between ontology topics and not necessarily between people.

3) *Strategy 3: Hybrid*: This strategy is a mix of the two former strategies: People are asked to be grouped together by affinity and people with same affinity will have similar (but not quite the same) set of pictures.

Alternatively, they can be asked to choose groups of two pictures among $\{i, j\}$, $\{l, m\}$, $\{r, s\}$, $\{u, v\}$ and $\{x, y\}$ (see Figure 2), but the former option is better if feasible.

Participants will be grouped by affinity such that, the closer the subjects, the larger number of pictures they share. This aims at counterbalancing the effect of pictures: indeed, if non close subjects have the same pictures, they may end up with close ontologies anyway.

- pros: Since topics and affinity are set together, there should at least be similarity to be found.
- cons: There is still a part of the correlation that will be measured by the result which will be due to the topics.

4) *Strategy 4: Indirect affinity grouping*: Each subject is given a set of 10 pictures. Participants are allowed to trade their pictures (exchange pictures one-by-one).

This can be absolute: The more picture in common, the highest the affinity.

This may be relative: If A has traded for x against y to B , and C has x . Then A is closer to C than to B . Moreover, if C has obtained x by trading it with D , then A is also closer from C than D .

- pros: Topics will induce affinity. The organisation is obtained spontaneously. This reproduces basically things that happen in social (exchange) networks.
- cons: It is not possible to ensure that this method will induce a clear order without too many pictures. Order agreement is difficult to test.

C. Population

We can consider three kind of audience to run the experiment:

- our teams
 - pros: People are available.
 - cons: They are not really available. They may be biased (at least this can be argued).
- students
 - pros: This is the usual procedure, which allows to run the experiment quickly.
 - cons: They should be sufficiently trained.
- outsiders: exploiting a structure like a VoCamp may be an idea.
 - pros: Enthousiasm.
 - cons: There is no control over the real population.

D. Measuring agreement

In order to measure the agreement between ontology similarity and subject affinity there may be different modalities. Indeed, no one expect that similarity and affinity will be expressed on the same scale and coincide exactly. What is important is the order induced by such affinities.

Hence, the agreement may be tested through different strategies.

1) *Strategy 1: Order agreement*: It is possible to compare the two orders and compute an agreement index. For this purpose, there exists measures such as Kendall's rank correlation and its variants [10]. They have been adapted in order to consider partial orders.

- pros: Easy to evaluate.
- cons: Require full order data from all subjects.

2) *Strategy 2: Cluster agreement*: Another opportunity is to use affinity and similarity for defining clusters in the graphs of subjects (which should reveal communities) and comparing such clusters.

- pros: Closer to the particular task of community mining.
- cons: Depends on the clustering technique used.

E. Specific problem

There is a specific problem related to the evaluation of the hypothesis. Indeed, it is claimed that:

- 1) distance between ontologies is a good approximation of distance between people;
- 2) distance between ontologies may reveal unknown proximity between people.

What can happen with such an experiment:

- we find that distance corresponds to affinity: this is fine, this validates claim (1);
- we find that distance does not correspond to affinity: we still can pretend that this validates claim (2).

So the results of such experiments must be taken with care. For instance, it could be possible that within each identified group of people we validate claim (1) and across groups, we

validate claim (2). This would have to be achieved by checking a posteriori the validity of (2).

In the remainder, we describe an actual experiment for testing the hypothesis.

IV. EXPERIMENT EXECUTION

The goal of this experiment is to obtain a semantic social network and affinity between people involved in this network.

We first tried to implement Strategy 4 of network acquisition and a variant of Strategy 2 of affinity acquisition (through Flickr), but we failed to gather sufficient information about affinity. Then we resorted to use Strategy 5 for network acquisition and the Strategy 1 for affinity acquisition on some of our students. For that purpose we needed to have people creating ontologies. In order to stimulate them, we presented this task as an exam in ontology engineering.

We also needed that people reveal their proximity with each others. However, it is quite difficult to obtain proximity assessment and ontologies in a population we had to resort to questionnaires.

A. Material and methods

1) *Population*: The subjects were $|U| = 19$ undergraduate students of a university taking an AI course. They were given one hour training in ontologies and Protégé.

The task was given to the students as a rated homework to be returned within one week. The subjects were asked to complete the task individually.

After all homeworks were completed and rated, subjects were asked to fill a questionnaire about their opinion about various aspects of the task.

2) *Task*: Participants have been given $|P| = 22$ pictures each representing some wildlife scene featuring an animal. They have been asked, as homework, to create an ontology and to annotate the pictures. They have been told that the marks would be on the correctness of ontologies and of the design of ontologies and annotations.

Each subject has been asked to return one week later an ontology and annotations of these pictures using the ontology.

See Appendix¹ for the instructions given to students.

In addition, students have been asked, through a questionnaire (Appendix ??), to provide information about their affinity with other students. In order to not introduce a too large bias, this questionnaire was about different aspects of the task, e.g., time spent, strategy. Affinity information was collected through direct assessment (Which other students do you feel have the same state of mind as yours) as well as indirect information (What is your favorite animal). We only exploited the former.

3) *Measure*: From the ontologies provided by the subjects, we computed similarity between subjects with similarity measures provided by the OntoSim library. OntoSim is an API dedicated to the computation of distance or similarity measures between ontologies and ontology entities. It provides measures used for matching ontologies and support for the development

of new measures. We tested both ontology space and alignment space measures [11].

More precisely we considered:

- a vector space measure based on the cosine distance in a multidimensional space where each ontology corresponds to a point represented by the TFIDF measure on the terms used in the ontology;
- an ontology space distance using a maximal coupling of the ontology entities based on their names;
- an ontology space distance using a maximal coupling of the ontology entities based on the triples they appear in;
- two alignment space measures based on the paths between ontologies;
- two alignment space measures based on the coverage of the alignment paths between ontologies;

For alignment space measures, we computed all alignments between pairs of ontologies based on the edit-distance on class names and we applied a threshold of .7 to the resulting alignment.

B. Collected ontologies

In Appendix² we display the class identifiers used by different users. In **green** we display those terms which are used by at least one other user. Terms followed by an asterisk (*) need correction before being syntactically the same. In **orange**, are those terms that can easily be identified as a hyponym/hypernym of another mostly because there is a restricting adjective. In **blue**, there is a problem that will be explained latter.

Among the 19 students, 4 returned ontologies without identified classes and have been discarded. We have tried to further group the terms found in ontologies in relatively natural domains (that may also be called modules). We identified six of them: species as found in the classification of species, common names used for animals, area location which have been massively used and locomotion means, diet and gender which can be identified but have been less used. Table³ provides the use of these different groups by the students. We only counted true classes and ignored those classes which identified ontologies like SexValue, habitat or diet.

Table⁴ presents the various parameters of the domains as they were used by the subjects. It shows how domains can connect users.

1) *Analysis*: We can draw several lessons from this set of data.

The first remark is that there is a lot of connections between these ontologies. We already have a network of naive ontologies which are heavily naturally interconnected. These connections come both from students using the same standardised module (species) and from students using the same common names (when the need is the same and the population homogeneous, there are many trivial matches). This

¹<http://intelligent.pe.kr/papers/2012/CCI-SMC/Appendix/Appendix.pdf>

²<http://intelligent.pe.kr/papers/2012/CCI-SMC/Appendix/Appendix.pdf>

³<http://intelligent.pe.kr/papers/2012/CCI-SMC/Tables/Table1.pdf>

⁴<http://intelligent.pe.kr/papers/2012/CCI-SMC/Tables/Table2.pdf>

remark seems to be important because currently we are lacking examples of networks of ontologies but it seems that they are not that difficult to produce.

The second remark is that there is an important presence of the identified domains: only 6 domains among 14 students, this is impressive, but the task was easy. These domains would certainly be the things on which people spontaneously agree. Moreover, in a task like this one, they converged towards the same kind of domains.

A third remark is the importance of standard vocabularies. Even if students were taught about reusing existing vocabularies, the importance of the reuse of the classification of species is stunning. This classification (whose terms may be easily found from wikipedia) provides many links between people. On the other side, many students felt the need to use some geographic ontology but they were not aware of any standard classification so they invented one, or rather they invented 6 of them, which do not match much.

The blue carnivora is a special case: there seems to be some students wanting to classify animals from their diet who used the carnivora concepts. But there are also students who found the Carnivora order in the classification of species and who found natural to add Herbivora as a taxon. However, there is no such class in the classification of species: Carnivora is a Mammalia characterised by the shape of the skull and dentition (which, of course may be related to eating meat) but does not refer to any diet (some Carnivora are not carnivore, and many carnivore are not Carnivora, like squales which are not mammals).

We only considered here the class labels in the ontologies. It would be interesting to consider the topology of the ontology. One conjecture is that those who have used the classification of species have also preserved its hierarchy.

C. Collected affinity

We have surveyed the students to know if there is some affinity between some of them. This allows for comparing their asserted affinity with the distances between their ontologies. This should also provide us with information about the likelihood that they shared their information.

1) *Results*: The results of the questionnaire are given in Table⁵. Students provided between one to three other students as affinity bearers.

2) *Analysis*: It is possible to cluster the directed graph provided by these assertions into 8 clusters of complete sub-graphs or 7 strongly connected components (see Figure⁶). Such clusters clearly represent affinity nests in the perception of subjects.

These are made of very few elements, thus we can conjecture that there are really strong affinity links (because affinity statements have been provided in both directions). Affinity outside these clusters may be considered secondary.

⁵<http://intelligent.pe.kr/papers/2012/CCI-SMC/Tables/Table3.pdf>

⁶<http://intelligent.pe.kr/papers/2012/CCI-SMC/Figures/Figure3.pdf>

V. AGREEMENT BETWEEN SIMILARITY MEASURES AND AFFINITY

Our goal is to check if there is some correlation between the affinities given by students and those inferred by ontology measures.

A. Method

For each selected measure, we computed the similarity value between each pair of ontologies. From these measures, we established a ranking from the closest ontology to the farthest one. Then, for each ontology, we compared the average ranking of ontologies used by the subjects selected by the student with the best ranking we can obtain with the same number of ontologies.

Given \vec{u} the ranks (in ascending order) of the n ontologies used by subjects with which a student has declared affinity and \vec{v} all the ranks (in ascending order), the agreement between the similarity and the the student affinities is defined as follow:

$$agreement(\vec{u}, \vec{v}) = 1 - \frac{\sum_{i=1}^n (u_i - v_i)}{n}$$

We also have considered two different cluster comparisons:

- the first one using hierarchical clustering of subjects based on their similarity;
- the second by defining the graph whose edges exists when the ontologies have a similarity higher than a threshold (here .4). This is similar to the graph of Figure⁷ and the same operation (component decomposition can be applied).

B. Result analysis

Table⁸ gives the agreements for each measure and each subject.

We reduced the alignment space path based measure to a single column since the results were the same (only the presence of an alignment path is taken into account and we had such alignments). These measures are not meaningful because all ontologies have the same rank. This remark also concerns VS(cosine+TFIDF) to a lower extent.

Alignment space measures more agreed with the given affinities than ontology space measures. This can be explained by the fact than some ontologies (2149, 2132, 1530 and 1643) do not contain named classes and then they are not comparable with other ontologies by using lexical features (used by these ontology space measures).

However, the level of correlation between affinity asserted by students and the one predicted by the measures are not particularly high, even at .72. This is confirmed by the two cluster analyses that we performed: the cluster computed by either hierarchical clustering or component analysis are largely different from those provided by affinity. Figure 3 displays the graph made of the highest union path coverage similarity for each subject (similarities within 3% are considered ties).

⁷<http://intelligent.pe.kr/papers/2012/CCI-SMC/Figures/Figure3.pdf>

⁸<http://intelligent.pe.kr/papers/2012/CCI-SMC/Tables/Table4.pdf>




Fig. 3: Neighborhood social graphs and its complete subgraphs obtained from ontology similarity (in bold are arrows representing remaining strongly connected components).

In summary, it seems that this experiment contradicts the hypothesis that similarity between ontologies reflects affinity between users in semantic social networks.

VI. CONCLUSION

The goal in this paper was to consider the experimental evaluation of hypotheses and measures for assessing people affinity in semantic social networks through ontology similarity. We offered various alternative settings for such an experiment and analysed the pros and cons of each solution. We performed one such experiment taking into account material availability. The results of this experiment are leaning towards invalidating the proposed hypothesis.

Several parameters may be considered that once corrected may lead to reconsider or to particularize these results: the maturity and training of students, the direct acquisition of affinity statement, or the fact that similarity indeed reveals unknown affinity (see §III-E) which cannot be ruled out.

It would be interesting to consider other suggested experimental settings to see if they confirm this result. Especially using real massive data sets would be an important step.

ACKNOWLEDGMENT

This work was supported by the Korea Foundation for International Cooperation of Science & Technology(KICOS) through a grant provided by the Korean Ministry of Education, Science & Technology (MEST) in 2009 (No. K20904000009).

REFERENCES

- [1] P. Mika, "Ontologies are us: A unified model of social networks and semantics," *Journal of web semantics*, vol. 5, no. 1, pp. 5–15, 2007.
- [2] J. Jung and J. Euzenat, "Towards semantic social networks," in *Proc. 4th European Semantic Web Conference, Innsbruck (AT)*, ser. Lecture Notes in Computer Science, vol. 4519, 2007, pp. 267–280.
- [3] P. Mika, "Flink: Semantic web technology for the extraction and analysis of social networks," *Journal of web semantics*, vol. 3, no. 2-3, pp. 211–223, 2005.
- [4] P. Haase, M. Ehrig, A. Hotho, and B. Schnizler, "Personalized information access in a bibliographic peer-to-peer system," in *Peer-to-Peer and Semantic Web*, S. Staab and H. Stuckenschmidt, Eds. Springer, 2006, pp. 143–158.
- [5] J. Yingzi, Y. Matsuo, and M. Ishizuka, "Extracting social networks among various entities on the web," in *Proc. 4th European Semantic Web Conference, Innsbruck (AT)*, ser. Lecture Notes in Computer Science, E. Franconi, M. Kifer, and W. May, Eds., vol. 4519. Springer, 2007, pp. 251–266.
- [6] Y. Matsuo, J. Mori, M. Hamasaki, T. Nishimura, H. Takeda, K. Hasida, and M. Ishizuka, "Polyphonet: An advanced social network extraction system from the web," *Journal of web semantics*, vol. 5, no. 4, pp. 262–278, 2007.
- [7] C. Tang, Z. Xu, and S. Dwarkadas, "Peer-to-peer information retrieval using self-organizing semantic overlay networks," in *Proc. ACM SIGCOMM 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, 2003.
- [8] A. Crespo and H. Garcia-Molina, "Semantic overlay networks for p2p systems," in *Proc. 3rd International Workshop on Agents and Peer-to-Peer Computing (AP2PC 2004)*, ser. Lecture Notes in Computer Science, vol. 3601. Springer, 2004, pp. 1–13.
- [9] M. d'Aquin and H. Lewen, "Cupboard – a place to expose your ontologies to applications and the community," in *Proc. 6th European Semantic Web Conference (ESWC), Heraklion (GR)*, ser. Lecture Notes in Computer Science, vol. 5554, 2009, pp. 913–918.
- [10] M. Kendall, *Rank correlation methods*. Griffin, London (UK), 1970.
- [11] J. David, J. Euzenat, and O. Sváb-Zamazal, "Ontology similarity in the alignment space," in *Proc. 9th international semantic web conference (ISWC), Shanghai (CN)*, 2010, pp. 129–144. [Online]. Available: <ftp://ftp.inrialpes.fr/pub/exmo/publications/david2010b.pdf>