Aligning Geospatial Ontologies Logically

Heshan Du

The University of Nottingham, UK

Abstract. Information sharing and updates have become increasingly important in the rapidly changing world. However, owing to the distributed and decentralized nature of information collection and storage, it is not easy to use information from different sources synergistically. Ontology plays as an important role in establishing formal descriptions of a domain of discourse. In geographic information science, the rapid developments of crowd-sourced geospatial databases challenge and also bring opportunities to the current geospatial information development framework. In this paper, a new semi-automatic method is proposed to align disparate geospatial ontologies, based on description logic and domain experts' knowledge.

1 Introduction

Information sharing and updates have become increasingly important in the rapidly changing world, with a large amount of disparate information available. However, owing to the distributed and decentralized nature of information development, it is not easy to fully capture the information content in different sources. A same expression can have different meanings in different context, and different expressions may refer to the same meaning. Such issues are quite common when disparate and related information sources exchange their data.

Ontology plays an important role in information sharing. Ontology, originated in the work of Aristotle, is a branch of philosophy which studies the existence of entities [41]. In computer science, ontology is an explicit formal specification of a shared conceptualization [14]. Compared to its origin, computer science ontology is not only about existence, but also about meaning, and making meanings as clear as possible [41]. Ontologies are often employed as important means for establishing explicit formal vocabulary shared among applications.

Description logics are a family of formalisms for representing the knowledge of an application domain [2]. They firstly define the relevant concepts (terminology), and then use these concepts to specify properties of objects in the domain [2]. Compared to traditional first-order logic, description logics provide tools for the manipulations of compound and complex concepts [7]. Description logics support inference patterns, including classification of concepts and individuals, which are often used by humans to structure and understand the world [2]. When dealing with ontology issues, description logics can be used as the logical underpinning.

The rapid developments in geographic information science emphasize the importance of geospatial information sharing. Spatial Data Infrastructures (SDI) refers to an institutional and technical foundation of policies, standards and procedures that enable organizations at multiple levels and scales to discover, evaluate, use and maintain

R.K. Rendsvig and S. Katrenko (Eds.): *ESSLLI 2012 Student Session Proceedings*, CEUR Workshop Proceedings vol.: see http://ceur-ws.org/, 2012, pp. 1–12.

geospatial data [28]. Over the last few years, the current top-down approach to SDI has been challenged by the rapid pace of technological development [17]. There is a need to address the separation of national and international SDI from crowd-sourced geospatial databases [1]. Relying on volunteers for data collection, crowd-sourced data is less expensive than authenticated data. In addition, although typically not as complete in its coverage or as consistent in its geometric or metadata quality as authenticated data, crowd-sourced data may provide a rich source of complementary information with the benefit of often more recent and frequent updates than that of authenticated data [18]. It is desirable to use authenticated and crowd-sourced geospatial data synergistically.

Compared to other ontologies, geospatial ontologies have several special properties. Firstly, many words within geospatial ontologies are often more widely used in daily life, and there is less consensus about their definitions. For example, the word 'creek' can refer to a river in Australia, while cannot in the US. The word 'field' has different meanings (e.g. a branch of knowledge, a piece of land, etc.) for different people in different contexts. There are no precise formal definitions which can tell 'river' and 'stream', 'lake' and 'pond' apart. In addition, geospatial ontologies often do not have a huge number of classes as biomedical or bioinformatics ontologies do, but may have many instances, referring to real world objects, whose locations, at least in theory, are verifiable. With respect to these properties and the underspecification of geospatial ontologies, no fully automated system can ensure the correctness and completeness of generated mappings. Therefore, experts are inevitably needed to make decisions, for example on the correctness of correspondences, based on their domain knowledge, which is often implicit in the individual ontologies. This research will finally lead to the minimisation of the human intervention.

This project aims to explore logic-based approaches to aligning disparate and related geospatial ontologies to obtain harmonized and maximized information content. Ontologies are disparate if they are created independently. Ontologies are related if they contain more than one common concept. Aligning means establishing relations between the vocabularies of two ontologies [9]. The desired output will be a collection of verified such relations, for query answering over multiple ontologies. If ontologies cannot be aligned logically, a deficiency report will be produced explaining reasons and providing suggestions about further actions to take in order to align them. In this paper, we propose a new semi-automatic method to align disparate geospatial ontologies, based on description logic and domain experts' knowledge. It is assumed that original information within ontologies is believed all the time as premises, whilst generated information, including disjointness axioms and mappings, is believed by default as assumptions, which may be retracted later. Differing from other existing methods, generated disjointness axioms are seen as assumptions, which are retractable during the overall aligning process. Disparate ontologies are aligned by finding a coherent and consistent assumption set with respect to them. Based on this main idea, algorithms are designed to align ontologies at the terminology level and the instance level. With respect to the special properties of geospatial ontologies, an algorithm is designed for refining correspondences between geospatial individuals taking their geometries and semantics into account.

The rest of the paper is organized as follows. The related work is summarized in Section 2. Section 3 introduces the geospatial ontologies we use, and explains some results generated by a state-of-the-art system called S-Match [12]. Our method is discussed in Section 4. Finally, it provides conclusions in Section 5.

2 Related Work

Ontology matching is the task of finding correspondences between entities from different ontologies [11]. A correspondence represents a semantic relation, such as inclusion or equivalence. A mapping is defined as a set of these correspondences [25]. Many ontology matching methods and systems have been proposed and developed in recent years [11] [35], based on shared upper ontologies, if available, or using other kinds of information, such as lexical and structural information, user input, external resources and prior matches [30]. The existing methods can be classified into three broad categories [4]. *Syntactic methods* rely on a syntactic analysis of the linguistic expressions to generate mappings. Though these methods are direct and effective, semantic relations between entities cannot be captured. *Pragmatic methods* infer the semantic relations between concepts from associated instance data. Though they work well when the instance data is representative and overlapping, this kind of methods use a strong form of induction, thus lack correctness and completeness. *Conceptual methods* compare the lexical representation of concepts to compute mappings. Socially negotiated meanings (e.g. dictionaries) are often used when generating relations, making the problem very complicated.

Many methods are hybrid, combining different approaches and making use of structural, lexical, domain or instance-based knowledge. Most of them apply heuristics-based or machine learning techniques to various characteristics of ontology. However, mappings generated by these methods often contain logical contradictions. Some systems, such as CtxMatch [5] and its extension S-Match [12], and more recently, Automated Semantic Mapping of Ontologies with Validation (ASMOV) [20], Knowledge Organization System Implicit Mapping (KOSIMap) [33], logiC-based ONtology inTEgratioN Tool using MAPpings (ContentMap) [22], LogMap [21], and Combinational Optimization for Data Integration (CODI) [29], seem to be exceptions, since they employ logical reasoning for either mapping generation or verification. Due to limited space, not all of them are discussed in detail.

CtxMatch and its extension S-Match are early logic-based attempts for ontology matching. While most of the other methods compute linguistic or structural similarities, CtxMatch shifts the semantic alignment problem to deducing relations between logical formulae [5, 6]. Relevant lexical (concepts denoted by words), world (relations between concepts) and structural knowledge is represented as logical formulae, and logical reasoning is employed to infer different kinds of semantic relations [34]. WordNet [27], an external resource, is employed to provide both lexical and world knowledge. The S-Match system re-implements and extends CtxMatch. Taking two tree-like structures (e.g. hierarchies) as input, it computes the strongest semantic relations between every pair of concepts [12]. The semantic matching has two main steps. Firstly, at element level, relations between labels are calculated, and then, at structure level, it generates relations between concepts, whose semantics are constructed based on the semantics of labels [36]. The structure level matching task is converted into propositional validity problems, and the standard DPLL-based SAT solver [3] is employed to check the unsatisfiability of propositional formulae [13]. However, S-Match only uses information in the tree-like structures extracted from ontologies, which is insufficient to guarantee the overall coherence of ontologies after applying the mapping relations. More recently, some matching tools have been developed, involving semantic verification into the alignment process.

LogMap [21] is a logic-based and scalable ontology matching tool. It addresses the challenges when dealing with large-scale bio-medical ontologies with tens (even hundreds) of thousands of classes. It employs lexical and structural methods to compute an initial set of mapping relations as the starting point for further discovery of mapping relations. The core of LogMap is an iterative process which alternates repair and discovery steps. In the repair step, unsatisfiable classes will be detected using propositional Horn representation and satisfiability checking, and be repaired using a greedy diagnosis algorithm. However, the propositional Horn satisfiability checking is sound but incomplete, and the underlying semantics is restricted to propositional logic, and thus cannot guarantee the coherence of the mapping between more expressive ontologies. In the discovery step, new mapping relations will be generated based on the similarity between classes which are semantically related to matched classes. ISUB [37] is employed to compute the similarity scores. Mapping relations which are newly discovered are active, and only active mapping relations can be eliminated in the repair step, whilst mapping relations found in earlier iterations are seen as established or valid. In other words, each mapping relation will be checked once, against the available information at that time, which, however, cannot guarantee its correctness when new information is discovered later.

Combinational Optimization for Data Integration (CODI) [29] is a probabilistic logical alignment system. It is based on Markov logic [10], which combines first-order logic with undirected probabilistic graphical models. As the main advantages over other existing matching approaches, Markov logic can combine hard logical axioms and soft uncertain formulae for potential correspondences. Cardinality constraints, coherence constraints and stability constraints are formalized using logical axioms and similarity measures. The matching problem is transformed to a maximum-a-posteriori optimization problem subject to these constraints. The GUROBI optimizer [15] is employed to solve the optimization problems. According to Noessner and Niepert [29], CODI reduces incoherence during the alignment process for the first time, compared to all other existing methods repairing alignments afterwards. CODI is based on the rationale of finding the most likely mapping by maximizing the sum of similarity-weighted probabilities for potential correspondences. It can be argued that during the optimization process, some valid correspondences can be thrown away. In addition, the input coherence constraints will influence the resulting mapping, however, in practice, many ontologies are underspecified, within which valid disjointness axioms are not always available.

In addition, there is some recent work on debugging and repairing ontologies and mappings in ontology networks [24] [32] [40] [23], which is still at an early stage. However, all of them use disjointness axioms as premises, rather than assumptions, and none of the matching systems discussed above have addressed the special properties of geospatial information fully. Several ontology-driven methods have been developed for integrating geospatial terminologies. Most of them are based on similarity measures or a predefined top-level ontology, and logical reasoning is only employed when formal ontologies commit to the same top-level ontology [8]. However, when ontologies are developed independently, the common top-level ontology is not always available. Additionally, there exist some other methods, such as [38] and [19], following the pragmatic approach to link geospatial schemas or ontologies, inferring the terminology correspondences from the instances correspondences. As discussed above, relying on a very strong form of induction from particular to the universal, this approach will lead to the lack of correctness and completeness [4]. Therefore, more research is required to fill in the gap, exploring logic-based approaches to aligning disparate geospatial ontologies.

3 Disparate Geospatial Ontologies

The Ordnance Survey of Great Britain (OSGB) ontology [16] and the OpenStreetMap (OSM) controlled vocabularies [31] are selected to undertake initial research. OSGB and OSM are representatives of authenticated and crowd-sourced geospatial information sources respectively. OSGB is the national topographic mapping agency of Great Britain. It has built ontologies for Hydrology and for Buildings and Places [16]. OSM is a collaborative project aimed to create a free editable map of the world [31]. It employs the bottom-up approach, relying on volunteers to collect the data. Currently, OSM does not have a standard ontology, but maintains a collection of commonly used tags for main map features [31]. An OSM feature ontology is generated automatically from the existing classification of main features. Both ontologies are written in the OWL 2 Web Ontology Language [39]. The OSGB Buildings and Places ontology has 692 classes and 1230 logical axioms, and its DL expressivity is $ALCHOIQ$. There are 663 classes and 677 logical axioms in the OSM ontology, whose DL expressivity is AL. Both ontologies, containing *no* disjointness axioms, are coherent.

To understand the ontologies more deeply, S-Match is employed to generate relations between concepts from them. To distinguish concepts from different ontologies, let us label each concept with the abbreviated name of the ontology it belongs to, such as $OSGB:School$ and $OSM:School$. All the labelled concepts will be treated as belonging to one super ontology. A relation then can be represented as an axiom within which all the concepts are labelled, like $OSGB:School \sqsupseteq OSM:School$.

Some of the mapping axioms generated by S-Match seem reasonable, such as $OSGB:Roof \equiv OSM:Roof$, $OSGB:Service \sqsupseteq OSM:Service$, and $OSGB:Accommodation \sqsupseteq OSM:Accommodation$. However, there are also some problematic relations. For example, $OSGB:Thing \equiv OSM:Nothing$ is derived, because the string 'Thing' is considered to be close enough by the string-based matcher [1] for stating that it is equivalent to the string 'Nothing'. The relation $OSGB:Person \sqsupseteq OSM:GuestHouse$ is generated because, 'GuestHouse' is split to 'Guest' and 'House', and 'House' is treated as a person's name, referring to a particular individual of 'Person'. The relation $OSGB:Person \sqsupseteq OSM:Dentist$ seems reasonable, just looking at it alone. However, $OSM:Dentist \sqsubseteq OSM:Healthcare$, and $OSM:Healthcare \sqsubseteq OSM:Amenity$. The $OSM:Dentist$ is used for tagging a place where a dentist practice or surgery is located, rather than referring to a person who is a dentist. In this case, S-Match seems too restrictive to deal with the informal use of terms and their variable meanings in crowd-sourced databases, such as OSM.

4 Method

Ontology alignment has attracted the attention of people working in several research fields, such as linguistics, philosophy, psychology and computer science. To fully understand and solve the problem, relying on only one approach is inadequate. Different approaches may play different important roles and solve different aspects of the problem effectively. This project focuses on the logic-based approach, and explores what logic can do and how far logic can go when aligning disparate geospatial ontologies.

[1] When matching labels of concepts, S-Match employs string-based, sense-based and gloss-based matchers [13].

To represent and reason with two ontologies O^i and O^j, where i, j are their names, as well as the matching relations between them, as if they all belong to one super ontology $O^i \cup O^j$, we label all atomic concepts and roles in each ontology by the name of the ontology. The ontology O^i is the set $\{\varphi^i : \varphi \in O^i\}$, where φ denotes a logical formula. A logical formula φ^i is labelled inductively as follows.

- $A^i = i : A$, for atomic concept A;
- $R^i = i : R$, for atomic role R;
- $(\neg B)^i = \neg B^i$;
- $(B \sqcap C)^i = B^i \sqcap C^i$;
- $(B \sqcup C)^i = B^i \sqcup C^i$;
- $\{o\}^i = \{o^i\}$, for nominal $\{o\}$, individual name o;
- $(\forall R.B)^i = \forall R^i.B^i$;
- $(\exists R.B)^i = \exists R^i.B^i$;
- $(\geq n\ R.C)^i = \geq n\ R^i.C^i$
- $(\leq n\ R.C)^i = \leq n\ R^i.C^i$
- $(= n\ R.C)^i == n\ R^i.C^i$
- $(R^-)^i = (R^i)^-$;
- $(R^+)^i = (R^i)^+$;
- $(B \sqsubseteq C)^i = B^i \sqsubseteq C^i$
- $(S \sqsubseteq T)^i = S^i \sqsubseteq T^i$

where B, C denote concept descriptions, S, T denote roles. Similarly, we label all individual names in each ontology by the ontology name:

- $a^i = i : a$;
- $(C(a))^i = C^i(a^i)$
- $(R(a, b))^i = R^i(a^i, b^i)$

where a, b denote individual names.

A terminology mapping is a set of correspondences between classes from different ontologies. A terminology correspondence is represented as one of the two basic forms:

$$B^i \sqsubseteq C^j \tag{1}$$

$$B^i \sqsupseteq C^j \tag{2}$$

where B, C denote class descriptions. The relation (1) states that the class B from the ontology i is more specific than or equivalent to the class C from the ontology j. The relation (2) states that the class B from the ontology i is more general than or equivalent to the class C from the ontology j. The equivalence relation (3) holds if and only if (1) and (2) both hold.

$$B^i \equiv C^j \tag{3}$$

It states that the concept B from the ontology i and the concept C from the ontology j are equivalent.

A disjointness axiom states that two or more classes are pairwise disjoint, having no common element. For example, $Person$ and $Place$ are disjoint. The disjointness axioms in ontologies play an important role in debugging ontology mappings. However, within the original geospatial ontologies, disjointness axioms are not always available or sufficient. Adding disjointness axioms manually, especially for large ontologies, is time-consuming and error-prone. Many existing systems employ more automatic approaches, either assuming the disjointness of siblings (e.g. KOSIMap [33]), or employing machine learning techniques to detect disjointness (e.g. [26]). After the disjointness axioms are generated by whatever means, all existing ontology matching or debugging methods, to the best of our knowledge, use them as premises, believing their validity through the overall following process, though the input disjointness axioms can be insufficient or too restrictive. Differing from these methods, we use generated disjointness axioms as assumptions, rather than premises, and ensure the assumption set is coherent. A disjointness assumption is represented as:

$$B \sqsubseteq \neg C \tag{4}$$

where B, C denote class descriptions either from ontology i or ontology j. We follow the terminology from [26], and adapt them to this context.

Definition 1 (Coherence). *An ontology O is coherent if there is no class C such that $O \models C \sqsubseteq \bot$. Otherwise, it is incoherent.*

Definition 2 (Coherence of an Assumption Set). *An assumption set A_s is incoherent with respect to an ontology O, if $O \cup A_s$ is incoherent, but O is coherent. Otherwise, it is coherent with respect to an ontology O.*

Definition 3 (Minimal Incoherent Assumption Set). *Given a set of assumptions A_s, a set $C \subseteq A_s$ is a minimal incoherent assumption set (MIA) iff C is incoherent and each $C^{'} \subset C$ is coherent.*

A minimal incoherent assumption set can be fixed by removing any axiom from it. When a MIA contains more than one element, one needs to decide which axiom to remove. Most of the existing methods remove the one either with the lowest confidence value or which is the least relevant. However, it can be argued that there is no consensus with respect to the measure of the degree of confidence or relevance. In addition, the confidence values or the relevance degrees might be unavailable, difficult to compute or compare. In such cases, it seems sensible to allow domain experts to make such decisions.

When aligning ontologies using a terminology mapping, Definition 2 is extended from one ontology O to two ontologies O_1 and O_2, given that the union of two ontologies $O_1 \cup O_2$ is an ontology. Based on these definitions, *Algorithm 1* is designed as follows [2]. An assumption in a minimal incoherent assumption set can be a disjointness axiom or a terminology correspondence axiom. The set of minimal incoherent assumption sets will be visualized clearly (Line 7). Domain experts are consulted to decide which assumption(s) to retract (Line 8). A repair action can be retracting or adding an assumption axiom. Users are allowed to take several repair actions at one time.

ALGORITHM 1: Terminology Level Alignment
Input: O_1, O_2: *coherent* ontologies
D_s: a disjointness assumption set
M_{st}: a terminology mapping between O_1 and O_2
Output: A_s: a coherent assumption set with respect to $O_1 \cup O_2$
1. $O := O_1 \cup O_2$
2. $A_s := D_s \cup M_{st}$
3. **assert** O is *coherent*

[2] In an algorithm, lines marked with $*$ may require manual intervention.

$O := O \cup A_s$
while O is *incoherent* **do**
  $S_{mia} := MIA(O)$
  $visualization(S_{mia})$
  $repair(O, S_{mia})$
  $update(A_s)$
end while
return A_s

Following the algorithm above, even if the problematic relations generated by S-Match are introduced, they can be retracted, given the disjointness assumptions such as $OSGB : Person \sqsubseteq \neg OSM : Accommodation$ and $OSGB : Person \sqsubseteq \neg OSM : Amenity$.

An instance level mapping is a set of individual correspondences. An individual correspondence is represented in one of the following forms:

$$(a^i, b^j) \in sameAs \tag{5}$$

$$(a^i, b^j) \in partOf \tag{6}$$

where a, b denote individual names. The relation (4) states that the individual name a from the ontology i and the individual name b from the ontology j refer to the same object. The relation (5) states that the individual name a from the ontology i refers to an object which is a part of the object the individual name b from the ontology j refers to.

When working with geospatial instances, *Algorithm 2* is designed to refine the initial instance level mapping using geometry, lexical and cardinality properties.

ALGORITHM 2: Refining GeoInstance Mapping
Input: M_{sa}: an initial instance level mapping for geospatial individuals
Output: M_{sa}: the refined input M_{sa}
1. **for each** individual correspondence m in M_{sa} **do**
2. $a_1 := m.individual_1$, $a_2 := m.individual_2$
3.   **if** $a_1.geometry, a_2.geometry$ are not *matched* **then**
4. $remove(M_{sa}, m)$
5.   **else if** $a_1.lexicons, a_2.lexicons$ are not *matched* **then**
6. $remove(M_{sa}, m)$
7.   **end if**
8. **end for**
9. **for each** individual b appearing more than once in M_{sa} **do**
10. $M_{s_b} := allCorrespondencesInvolving(b)$
11. $repair(M_{sa}, M_{s_b})$
12. **end for**
13. **return** M_{sa}

Given a set of correspondences linking geospatial individuals from different ontologies, *Algorithm 2* applies three main constraints, these are, geometry, lexical and cardinality. Firstly, a correspondence is invalid if the geometries of the linked individuals are not matched (Line 3-4). For example, when the geometries are both polygons, if they are spatially disjoint, they cannot be matched. Secondly, a correspondence is invalid if the lexicons, i.e. meaningful labels indicating identity, cannot be matched (Line 5-6).

The lexical matching required should be robust enough to tolerate partial differences in labelling. For example, a full name and its abbreviation should be matched. Thirdly, if an individual is involved in several different 'sameAs' correspondences, then these correspondences need to be repaired (Line 9-12), for example, by changing the relation from 'sameAs' to 'partOf'. The three constraints complement each other to cope with the following possibilities. Different geospatial individuals may share the same label or the same location in an ontology. In addition, the same geospatial individual may be represented as a whole in one ontology, whilst as several parts of it in the other.

The algorithm for aligning instances is generated by extending the assumption set to include instance correspondences (output of *Algorithm 2*) and changing coherence checking to consistency checking in *Algorithm 1*. Similarly, domain experts are consulted to make decisions to repair inconsistencies. Consider the example below.

Example 1. $OSGB : 1000002308476718$ refers to a $OSGB : HealthCentre$ labelled as 'SNEINTON HEALTH CENTRE'. $OSM : 62134030$ refers to a $OSM : Clinic$ labelled also as 'SNEINTON HEALTH CENTRE'. Their geometries are very similar. However, the existence of the following assumptions can lead to inconsistency.

$$(OSGB : 1000002308476718, OSM : 62134030) \in sameAs \tag{7}$$

$$OSGB : Clinic \equiv OSM : Clinic \tag{8}$$

$$OSGB : Clinic \sqsubseteq \neg OSGB : HealthCentre \tag{9}$$

Domain experts are consulted to decide which assumption(s) to retract. To keep the individual correspondence (7), it is reasonable to retract (9). This differs from all other methods, which use (9) as a premise, and therefore will either remove (7) or (8), though both are reasonable.

This method has been implemented as a system. Its performance is being evaluated, compared to other existing systems, such as S-Match, CODI and LogMap.

5 Conclusion

To facilitate the geospatial information sharing and updates, it is important to harmonize disparate and related geospatial ontologies. This paper discusses problems involved, and presents a new logic-based method to deal with them. Future work includes employing a truth maintenance system to track logical dependencies and qualitative spatial reasoning to check topological consistency of geospatial data.

References

1. Anand, S., Morley, J., Jiang, W., Du, H., Jackson, M.J., Hart, G.: When worlds collide: Combining Ordnance Survey and OpenStreetMap data. In: Association for Geographic Information (AGI) GeoCommunity '10 Conference (2010)
2. Baader, F., Calvanese, D., McGuinness, D.L., Nardi, D., Patel-Schneider, P.F. (eds.): The Description Logic Handbook. Cambridge University Press (2007)
3. Berre, D.L., Parrain, A.: The Sat4j Library, release 2.2. Journal on Satisfiability, Boolean Modeling and Computation 7(2-3), 59–64 (2010)
4. Bouquet, P.: Contexts and ontologies in schema matching. In: Context and Ontology Representation and Reasoning (2007), `http://ceur-ws.org/Vol-298/paper2.pdf`

5. Bouquet, P., Serafini, L., Zanobini, S.: Semantic Coordination: A New Approach and an Application. In: International Semantic Web Conference. pp. 130–145 (2003), http://disi.unitn.it/~bouquet/papers/ISWC2003-CtxMatch.pdf
6. Bouquet, P., Serafini, L., Zanobini, S.: Peer-to-Peer Semantic Coordination. Web Semantics: Science, Services and Agents on the World Wide Web 2(1) (2004), http://www.websemanticsjournal.org/index.php/ps/article/view/54
7. Brachman, R.J., Levesque, H.J.: Knowledge Representation and Reasoning. The Morgan Kaufmann Series in Artificial Intelligence, Morgan Kaufmann (2004), http://books.google.co.uk/books?id=OuPtLaA5QjoC
8. Buccella, A., Cechich, A., Fillottrani, P.: Ontology-Driven Geographic Information Integration: A Survey of Current Approaches. Computers and Geosciences 35(4), 710 – 723 (2009), http://www.sciencedirect.com/science/article/pii/S0098300408002021
9. Choi, N., Song, I.Y., Han, H.: A survey on ontology mapping. SIGMOD Record 35, 34–41 (September 2006), http://doi.acm.org/10.1145/1168092.1168097
10. Domingos, P., Lowd, D., Kok, S., Poon, H., Richardson, M., Singla, P.: Just Add Weights: Markov Logic for the Semantic Web. In: Uncertainty Reasoning for the Semantic Web I, ISWC International Workshops, URSW 2005-2007, Revised Selected and Invited Papers. pp. 1–25 (2008)
11. Euzenat, J., Shvaiko, P.: Ontology Matching. Springer-Verlag, Heidelberg (DE) (2007), http://www.springerlink.com/content/978-3-540-49611-3#section=273178&page=1
12. Giunchiglia, F., Shvaiko, P., Yatskevich, M.: S-Match: an Algorithm and an Implementation of Semantic Matching. In: European Semantic Web Conference (ESWS). pp. 61–75 (2004)
13. Giunchiglia, F., Yatskevich, M., Shvaiko, P.: Semantic Matching: Algorithms and Implementation. Journal on Data Semantics IX 9, 1–38 (2007)
14. Gruber, T.R.: A translation approach to portable ontology specifications. Knowledge Acquisition 5, 199–220 (1993), http://citeseer.ist.psu.edu/viewdoc/summary?doi=10.1.1.101.7493
15. Gurobi Optimization Inc.: Gurobi Optimizer Reference Manual. http://www.gurobi.com (2012), http://www.gurobi.com
16. Hart, G., Dolbear, C., Kovacs, K., Guy, A.: Ordnance Survey Ontologies. http://www.ordnancesurvey.co.uk/oswebsite/ontology (2008)
17. Jackson, M.J.: The Impact of Open Data, Open Source Software and Open Standards on the Evolution of National SDI's. In: Third Open Source GIS Conference (21-22 June 2011), http://uiwapmds01.nottingham.ac.uk/qcsplace/ondemand/events11/a5a1e446858f4867b77716111a/player.HTM
18. Jackson, M.J., Rahemtulla, H., Morley, J.: The Synergistic Use of Authenticated and Crowd-Sourced Data for Emergency Response. In: 2nd International Workshop on Validation of Geo-Information Products for Crisis Management (VALgEO). Ispra, Italy (11-13 October 2010), http://globesec.jrc.ec.europa.eu/workshops/valgeo-2010/proceedings
19. Jain, P., Hitzler, P., Sheth, A.P., Verma, K., Yeh, P.Z.: Ontology Alignment for Linked Open Data. In: International Semantic Web Conference (1). pp. 402–417 (2010)
20. Jean-Mary, Y.R., Shironoshita, E.P., Kabuka, M.R.: ASMOV: results for OAEI 2010. In: the 5th International Workshop on Ontology Matching (OM-2010) (2010)
21. Jiménez-Ruiz, E., Grau, B.C.: LogMap: Logic-Based and Scalable Ontology Matching. In: International Semantic Web Conference (1). pp. 273–288 (2011)

22. Jiménez-Ruiz, E., Grau, B.C., Horrocks, I., Llavori, R.B.: Ontology Integration Using Mappings: Towards Getting the Right Logical Consequences. In: The 6th European Semantic Web Conference (ESWC). pp. 173–187 (2009)
23. Lambrix, P., Liu, Q.: Debugging is-a Structure in Networked Taxonomies. In: the 4th International Workshop on Semantic Web Applications and Tools for the Life Sciences. pp. 58–65. SWAT4LS '11, ACM, New York, NY, USA (2011), http://doi.acm.org/10.1145/2166896.2166914,http://www.informatik.uni-trier.de/~ley/db/indices/a-tree/l/Lambrix:Patrick.html
24. Meilicke, C., Stuckenschmidt, H.: An Efficient Method for Computing Alignment Diagnoses. In: Third International Conference on Web Reasoning and Rule Systems. pp. 182–196 (2009)
25. Meilicke, C., Stuckenschmidt, H., Tamilin, A.: Reasoning Support for Mapping Revision. Journal of Logic and Computation 19(5), 807–829 (2008), http://disi.unitn.it/~p2p/RelatedWork/Matching/Meilicke08reasoning.pdf
26. Meilicke, C., Völker, J., Stuckenschmidt, H.: Learning Disjointness for Debugging Mappings between Lightweight Ontologies. In: Proceedings of the 16th international conference on Knowledge Engineering: Practice and Patterns. pp. 93–108. EKAW '08, Springer-Verlag, Berlin, Heidelberg (2008), http://dx.doi.org/10.1007/978-3-540-87696-0_11
27. Miller, G.A.: WordNet: a lexical database for English. Communications of the ACM 38, 39–41 (November 1995), http://dl.acm.org/citation.cfm?doid=219717.219748
28. Nebert, D.D.: Developing Spatial Data Infrastructures: The SDI Cookbook. Global Spatial Data Infrastructure Association (GSDI) (2004)
29. Niepert, M., Meilicke, C., Stuckenschmidt, H.: A Probabilistic-Logical Framework for Ontology Matching. In: American Association for Artificial Intelligence (2010)
30. Noy, N., Stuckenschmidt, H.: Ontology alignment: An annotated bibliography. In: Semantic Interoperability and Integration. Schloss Dagstuhl (2005), http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.93.2574
31. OpenStreetMap: The Free Wiki World Map. http://www.openstreetmap.org (2012)
32. Qi, G., Ji, Q., Haase, P.: A Conflict-Based Operator for Mapping Revision: Theory and Implementation. In: Proceedings of the 8th International Semantic Web Conference. pp. 521–536. ISWC '09, Springer-Verlag, Berlin, Heidelberg (2009), http://dx.doi.org/10.1007/978-3-642-04930-9_33,http://www.springerlink.com/content/0v25w11500867103/fulltext.pdf
33. Reul, Q., Pan, J.Z.: KOSIMap: Use of Description Logic Reasoning to Align Heterogeneous Ontologies. In: the 23rd International Workshop on Description Logics (DL 2010) (2010)
34. Serafini, L., Zanobini, S., Sceffer, S., Bouquet, P.: Matching Hierarchical Classifications with Attributes. In: European Semantic Web Conference (ESWC). pp. 4–18 (2006)
35. Shvaiko, P., Euzenat, J.: Ontology Matching: State of the Art and Future Challenges. IEEE Transactions on Knowledge and Data Engineering (2012)
36. Shvaiko, P., Giunchiglia, F., Yatskevich, M.: Semantic Matching with S-Match, vol. Part 2, pp. 183–202 (2010)
37. Stoilos, G., Stamou, G.B., Kollias, S.D.: A String Metric for Ontology Alignment. In: International Semantic Web Conference. pp. 624–637 (2005)
38. Volz, S., Walter, V.: Linking Different Geospatial Databases by Explicit Relations. In: Proceedings of the XX th International Society for Photogramme-

try and Remote Sensing (ISPRS) Congress, Comm. IV. pp. 152–157 (2004), http://www.isprs.org/proceedings/XXXV/congress/comm4/papers/332.pdf
39. W3C: OWL 2 Web Ontology Language. http://www.w3.org/TR/owl2-overview (2009)
40. Wang, P., Xu, B.: Debugging Ontology Mappings: A Static Approach. Computers and Artificial Intelligence 27(1), 21–36 (2008), http://disi.unitn.it/~p2p/RelatedWork/Matching/MappingDebug_CAI.pdf
41. Welty, C.A.: Ontology research. AI Magazine 24(3), 11–12 (2003), http://www.aaai.org/ojs/index.php/aimagazine/article/view/1714