HAL
open science

A Formal Semantics for Weighted Ontology Mappings

Manuel Atencia, Alexander Borgida, Jérôme Euzenat, Chiara Ghidini, Luciano Serafini

To cite this version:

Manuel Atencia, Alexander Borgida, Jérôme Euzenat, Chiara Ghidini, Luciano Serafini. A Formal Semantics for Weighted Ontology Mappings. ISWC 2012 - International Semantic Web Conference, Nov 2012, Boston, MA, United States. pp.17-33, 10.1007/978-3-642-35176-1_2 . hal-00768413

HAL Id: hal-00768413

https://inria.hal.science/hal-00768413v1

Submitted on 21 Dec 2012

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Formal Semantics for Weighted Ontology Mappings

Manuel Atencia[1,2], Alexander Borgida[3], Jérôme Euzenat[1,2],
Chiara Ghidini[4], and Luciano Serafini[4]

[1] INRIA, France
{manuel.atencia,jerome.euzenat}@inria.fr
[2] University of Grenoble, France
[3] Rutgers University, United States
borgida@cs.rutgers.edu
[4] Fondazione Bruno Kessler, Italy
{ghidini,serafini}@fbk.eu

Abstract. Ontology mappings are often assigned a weight or confidence factor by matchers. Nonetheless, few semantic accounts have been given so far for such weights. This paper presents a formal semantics for weighted mappings between different ontologies. It is based on a classificational interpretation of mappings: if O_1 and O_2 are two ontologies used to classify a common set X, then mappings between O_1 and O_2 are interpreted to encode how elements of X classified in the concepts of O_1 are re-classified in the concepts of O_2, and weights are interpreted to measure how precise and complete re-classifications are. This semantics is justifiable by extensional practice of ontology matching. It is a conservative extension of a semantics of crisp mappings. The paper also includes properties that relate mapping entailment with description logic constructors.

1 Introduction

Ontology mappings are used to express semantic relations between components of two heterogeneous ontologies. They are key artifacts for the integration of knowledge encoded in distinct schemas. On the one hand, theoretical studies on ontology mappings give a formal background for *crisp mappings*, i.e. mappings that express set theoretical relations — subsumption ($\sqsubseteq$ and $\sqsupseteq$), equivalence ($\equiv$) and disjointness ($\perp$) — between the extensional meaning of the concepts and relations of two ontologies [3, 16]. On the other hand, the majority of the state-of-the-art tools that automatically match ontologies generate *weighted mappings*, i.e., crisp mappings associated with a confidence value, typically a real number between 0 and 1 [5]. There is, however, no shared view on how these weights should be interpreted, neither in ontology matching nor in related fields such as database schema matching [6].

In this paper we fill this gap by providing a novel formal semantics for interpreting the confidence value associated with a mapping. This semantics is not based on standard probabilistic notions such as ones used to extend Description Logics (DL) and rule-based descriptions of uncertainty, but instead is based on a classificational interpretation of mappings which reflects a family of approaches used in ontology matching techniques (see Chapter 4.4. in [5], and [15] as a recent example): if two ontologies

O_1 and O_2 are used to classify a common set X of items, mappings between O_1 and O_2 encode how elements of X classified in the concepts of O_1 are re-classified in the concepts of O_2, and weights measure how precise and complete the re-classifications are. We fall back on precision, recall, and F-measures, as they are used in the context of classification tasks, for the formalisation of weighted subsumptions and equivalence, respectively.

The proposed semantics makes it possible to discover inconsistencies and detect implications over sets of weighted mappings. In other words, to understand when a mapping can be derived from others, and when sets of mappings are inconsistent. We introduce a notion of logical consequence between weighted mappings, and investigate entailment between weighted mappings w.r.t. description logic constructors. Moreover, we prove that the semantics for weighted mappings introduced in the paper is a conservative extension of the semantics of crisp mappings for a specific class of Distributed Description Logics (DDLs) [3].

The paper is organised as follows. Section 2 introduces a formal semantics for weighted mappings, including the notion of (weighted) mapping entailment. In Sections 3–5 we show the adequacy of the proposed semantics, including the proof that it is a conservative extension of the semantics of DDL mappings; properties of the mapping entailment w.r.t. description logic constructors; and variants of the proposed semantics. In Section 6, we summarise related work, and we finish with concluding remarks.

2 Classificational Semantics for Weighted Mappings

The semantics presented in this paper is designed for weighted mappings between pairs of ontologies expressed in (a fragment of) first order language with tarskian semantics. However, for the sake of presentation we focus on a particular fragment of FOL, namely the description logic $\mathcal{ALCO}$: the basic DL logic $\mathcal{ALC}$ extended with nominals. We choose $\mathcal{ALCO}$ to present this semantics, since it is the smallest logic that contains full propositional connectives, relations and constants. $\mathcal{ALCO}$ is a DL logic defined on an alphabet $\Sigma = \mathcal{CN} \uplus \mathcal{RN} \uplus \mathcal{ON}$, where $\mathcal{CN}$ is a set of concept symbols, $\mathcal{RN}$ is a set of role symbols, and $\mathcal{ON}$ is a set of individual symbols. Complex concept expressions (simply called concepts) in $\mathcal{ALCO}$ are defined by the following grammar:

$$C, D := \top | \bot | A | \neg C | C \sqcap D | \exists R.C | \{o\}$$

where $A \in \mathcal{CN}$, $R \in \mathcal{RN}$ and $o \in \mathcal{ON}$. As usual, $\forall R.C$ stands for $\neg\exists R.\neg C$, and $C \sqcup D$ for $\neg(\neg C \sqcap \neg D)$. A general inclusion axiom (GCI) is an expression of the form $C \sqsubseteq D$ where C and D are concepts.

An interpretation $\mathcal{I}$ is a pair $\mathcal{I} = \langle \Delta^{\mathcal{I}}, \cdot^{\mathcal{I}} \rangle$ where $\Delta^{\mathcal{I}}$ is a non-empty set, called the domain of $\mathcal{I}$ and $\cdot^{\mathcal{I}}$ is a function such that $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$ for $A \in \mathcal{CN}$, $R^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$ for $R \in \mathcal{RN}$, and $o^{\mathcal{I}} \in \Delta^{\mathcal{I}}$ for $o \in \mathcal{ON}$. The interpretation of complex concepts of $\mathcal{ALCO}$ is defined according to the following rules:

$$\top^{\mathcal{I}} = \Delta^{\mathcal{I}} \qquad (C \sqcap D)^{\mathcal{I}} = C^{\mathcal{I}} \cap D^{\mathcal{I}}$$

$$\bot^{\mathcal{I}} = \emptyset \qquad \{a\}^{\mathcal{I}} = \{a^{\mathcal{I}}\}$$

$$(\neg C)^{\mathcal{I}} = \Delta^{\mathcal{I}} \setminus C^{\mathcal{I}} \qquad (\exists R.C)^{\mathcal{I}} = \{x \in \Delta^{\mathcal{I}} | \exists y, (x, y) \in R^{\mathcal{I}} \ \& \ y \in C^{\mathcal{I}}\}$$

An interpretation $\mathcal{I}$ satisfies a GCI $C \sqsubseteq D$, in symbols $\mathcal{I} \models C \sqsubseteq D$, if $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$. An ontology O is a set of GCIs. An interpretation satisfies O, in symbols $\mathcal{I} \models O$, if $\mathcal{I} \models \varphi$ for all $\varphi \in O$. A CGI is entailed by an ontology O, in symbols $O \models C \sqsubseteq D$ if $\mathcal{I} \models C \sqsubseteq D$, for all interpretations $\mathcal{I}$ that satisfy O.

$\mathcal{ALCO}$ has the "finite model property", meaning that if an entailment holds in all finite interpretations (ones with finite domain), then it is a theorem, in the sense that it holds for all interpretations. Not all DLs have the finite model property, e.g., adding cardinality constraints and inverses to $\mathcal{ALCO}$ prevents it.

2.1 Weighted Mappings

We start from the definition of mapping presented in [3] that we recall below.

Definition 1. *Let $\{O_i\}_{i \in I}$ be a family of ontologies. A* mapping *from O_i to O_j is an expression of the form*

$$i{:}C \; r \; j{:}D$$

where C and D are concepts of O_i and O_j, respectively, and $r \in \{\sqsubseteq, \equiv, \sqsupseteq, \bot\}$.

The notion of *weighted mapping* generalises it by associating mappings to a closed subinterval of $[0, 1]$.

Definition 2. *Let $\{O_i\}_{i \in I}$ be a family of ontologies. A* weighted mapping *from O_i to O_j is an expression of the form*

$$i{:}C \; r_{[a,b]} \; j{:}D$$

where C and D are concepts of O_i and O_j, respectively, $r \in \{\sqsubseteq, \equiv, \sqsupseteq, \bot\}$ and a, b are real numbers in the unit interval $[0, 1]$.

Remark 1. Notice that if $a > b$ then $[a, b] = \emptyset$ which is a closed subinterval of $[0, 1]$.

2.2 Formal Semantics for Weighted Mappings

Our semantics for weighted mappings is based on the following intuition. Assume that the concepts of two ontologies O_1 and O_2 are used to classify a common set of elements X. Mappings from O_1 to O_2 encode how the elements of X classified in the concepts of O_1 are re-classified in the concepts of O_2, and the weights encode how precise and complete these re-classifications are. Let us pin down this intuition and see how it can be used to define a formal semantics for a weighted mapping $1{:}C \; r_{[a,b]} \; 2{:}D$.

Let $X = \{x_1, \ldots, x_n\}$ be a non-empty finite set of fresh constants not occurring in $L(O_1)$ or $L(O_2)$.[5] The set X is meant to represent the set of shared items classified by concepts of the ontologies O_1 and O_2. A classification of X in O_1 is specified by virtue of an interpretation $\mathcal{I}_1$ of O_1 extended with the elements of X as follows. Let C

[5] Be aware that X is not a concept name, and that $\{x_1, \ldots, x_n\}$ is not a concept. Instead X is a meta-notation used in this paper to refer to an arbitrary finite set.

be a concept of O_1 and x_k a fresh constant of X; we say that x_k is classified under C according to $\mathcal{I}_1$ if $x_k^{\mathcal{I}_1} \in C^{\mathcal{I}_1}$. The set

$$C_X^{\mathcal{I}_1} = \{x \in X \mid x^{\mathcal{I}_1} \in C^{\mathcal{I}_1}\}$$

then represents the subset of items of X classified under C according to $\mathcal{I}_1$. Note that $C_X^{\mathcal{I}_1}$ is a subset of X whereas $C^{\mathcal{I}_1}$ is a subset of the domain of the interpretation $\mathcal{I}_1$. In addition, $C_X^{\mathcal{I}_1}$ is always a finite set while $C^{\mathcal{I}_1}$ may be infinite.

Let $\mathcal{I}_1$ and $\mathcal{I}_2$ be interpretations of O_1 and O_2, respectively, and let C and D be the concepts of O_1 and O_2, occurring in $1{:}C\ r_{[a,b]}\ 2{:}D$. Since we do not want to make any commitment on the interpretation domains of the two ontologies, it may be the case that the sets $C^{\mathcal{I}_1}$ and $D^{\mathcal{I}_2}$ cannot be compared as they might be defined over independent interpretation domains. Yet the sets $C_X^{\mathcal{I}_1}$ and $D_X^{\mathcal{I}_2}$ can be compared as they are both subsets of X which represent the sets of items of X classified under C according to $\mathcal{I}_1$ and under D according to $\mathcal{I}_2$, respectively. We can therefore examine the different types of mappings $1{:}C\ r_{[a,b]}\ 2{:}D$ obtained by looking at the different $r \in \{\sqsubseteq, \equiv, \sqsupseteq, \perp\}$.

Intuitively, the mapping $1{:}C \sqsubseteq 2{:}D$ is used to express that any item in X which is classified under C according to $\mathcal{I}_1$ is (re-)classified under D according to $\mathcal{I}_2$. The weighted mapping $1{:}C \sqsubseteq_{[a,b]} 2{:}D$ is thus used to express the fact that the proportion of items of X classified under C according to $\mathcal{I}_1$ which are (re-)classified under D according to $\mathcal{I}_2$ lies in the interval $[a, b]$. Assuming that $|C_X^{\mathcal{I}_1}| \neq \emptyset$, we can rewrite this intuition in the formula

$$\frac{|C_X^{\mathcal{I}_1} \cap D_X^{\mathcal{I}_2}|}{|C_X^{\mathcal{I}_1}|} \in [a, b] \tag{1}$$

which can be seen as the recall of $D_X^{\mathcal{I}_2}$ w.r.t. $C_X^{\mathcal{I}_1}$. Indeed, given two sets A and B, the *recall* of B w.r.t. A is defined as

$$R(A, B) = \frac{|A \cap B|}{|A|}$$

unless $|A| = 0$, in which case $R(A, B) = 1$. Thus, the condition in (1) can be rephrased as $R(C_X^{\mathcal{I}_1}, D_X^{\mathcal{I}_2}) \in [a, b]$.

Example 1. Let $X = \{x_1, \ldots, x_{10}\}$, and $C_X^{\mathcal{I}_1}$ and $D_X^{\mathcal{I}_2}$ as in the two following diagrams:

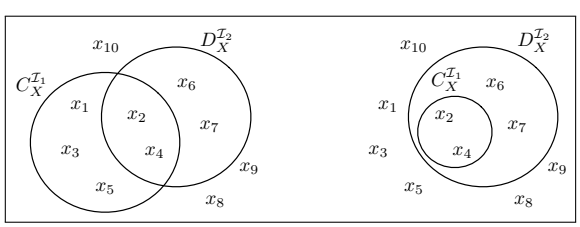

It is immediate to see that on the left hand side diagram, $R(C_X^{\mathcal{I}_1}, D_X^{\mathcal{I}_2}) = \frac{2}{5} = 0.4$ while on the right hand side diagram, $R(C_X^{\mathcal{I}_1}, D_X^{\mathcal{I}_2}) = \frac{2}{2} = 1$.

The weighted mapping $1\!:\!C \sqsupseteq_{[a,b]} 2\!:\!D$, in turn, is used to express the fact that the fraction of items of X classified by D according to $\mathcal{I}_2$ which are (re-)classified under C according to $\mathcal{I}_1$ lies in the interval $[a, b]$. Under the assumption that $|D_X^{\mathcal{I}_2}| \neq \emptyset$, we can rewrite this as:

$$\frac{|C_X^{\mathcal{I}_1} \cap D_X^{\mathcal{I}_2}|}{|D_X^{\mathcal{I}_2}|} \in [a, b] \tag{2}$$

which can be seen as the precision of $D_X^{\mathcal{I}_2}$ w.r.t. $C_X^{\mathcal{I}_1}$. The *precision* of B w.r.t. A is in fact given by

$$P(A, B) = \frac{|A \cap B|}{|B|}$$

unless $|B| = 0$, in which case $P(A, B) = 1$. Thus, the condition in (2) can be rephrased as $P(C_X^{\mathcal{I}_1}, D_X^{\mathcal{I}_2}) \in [a, b]$.

Example 2. Let $X = \{x_1, \ldots, x_{10}\}$, and $C_X^{\mathcal{I}_1}$ and $D_X^{\mathcal{I}_2}$ as in Example 1. It is immediate to see that in both cases $P(C_X^{\mathcal{I}_1}, D_X^{\mathcal{I}_2}) = \frac{2}{4} = 0.5$.

By keeping the parallelism with classification systems, the natural way to interpret the weighted mapping $1\!:\!C \equiv_{[a,b]} 2\!:\!D$ is by means of the F-measure, which is the harmonic mean of precision and recall. The *F-measure* of A and B is defined as

$$F(A, B) = 2 \cdot \frac{P(A, B) \cdot R(A, B)}{P(A, B) + R(A, B)}$$

unless $P(A, B)$ and $R(A, B)$ are equal to 0, then $F(A, B) = 0$. It can be expressed as

$$F(A, B) = 2 \cdot \frac{|A \cap B|}{|A| + |B|}$$

unless $|A| = |B| = 0$, in which case $F(A, B) = 1$. In this way, the weighted mapping $1\!:\!C \equiv_{[a,b]} 2\!:\!D$ encodes that $F(C_X^{\mathcal{I}_1}, D_X^{\mathcal{I}_2}) \in [a, b]$.

We conclude the above explanation with the definition of the degree of satisfiability of a mapping. Then we introduce mapping satisfiability and mapping entailment.

Definition 3 (Degree of satisfiability of a mapping). *Let O_i and O_j be two ontologies and let X be a non-empty finite set of fresh individual constants. Let $\mathcal{I}_i$ and $\mathcal{I}_j$ be two interpretations of O_i and O_j, respectively, extended with the set X. The* degree of satisfiability of the mapping $i\!:\!C\ r\ j\!:\!D$ with respect to the pair $(\mathcal{I}_i, \mathcal{I}_j)$ and X *is denoted by* $\mathrm{ds}_X(\mathcal{I}_i, \mathcal{I}_j, C, r, D)$ *and defined as*

$$\begin{aligned}
\mathrm{ds}_X(\mathcal{I}_i, \mathcal{I}_j, C, \sqsubseteq, D) &= R(C_X^{\mathcal{I}_i}, D_X^{\mathcal{I}_j}) \\
\mathrm{ds}_X(\mathcal{I}_i, \mathcal{I}_j, C, \sqsupseteq, D) &= P(C_X^{\mathcal{I}_i}, D_X^{\mathcal{I}_j}) \\
\mathrm{ds}_X(\mathcal{I}_i, \mathcal{I}_j, C, \equiv, D) &= F(C_X^{\mathcal{I}_i}, D_X^{\mathcal{I}_j}) \\
\mathrm{ds}_X(\mathcal{I}_i, \mathcal{I}_j, C, \bot, D) &= 1 - F(C_X^{\mathcal{I}_i}, D_X^{\mathcal{I}_j})
\end{aligned}$$

The pair $(\mathcal{I}_i, \mathcal{I}_j)$ satisfies the weighted mapping $i\!:\!C\ r_{[a,b]}\ j\!:\!D$ modulo X*, denoted by* $(\mathcal{I}_i, \mathcal{I}_j) \models_X i\!:\!C\ r_{[a,b]}\ j\!:\!D$*, if and only if* $\mathrm{ds}_X(\mathcal{I}_i, \mathcal{I}_j, C, r, D) \in [a, b]$*. If M is a set of weighted mappings from O_i to O_j, the pair* $(\mathcal{I}_i, \mathcal{I}_j)$ satisfies M modulo X*, in symbols,* $(\mathcal{I}_i, \mathcal{I}_j) \models_X M$*, if* $(\mathcal{I}_i, \mathcal{I}_j) \models_X m$ *for every* $m \in M$*.*

Definition 4 (Mapping satisfiability). *Let O_i and O_j be two ontologies and let X be a non-empty finite set of fresh individual constants. Let M be a set of weighted mappings from O_i to O_j. The set M is* satisfiable modulo X *if there exist interpretations $\mathcal{I}_i$ and $\mathcal{I}_j$ of O_i and O_j, respectively, such that $(\mathcal{I}_i, \mathcal{I}_j) \models_X M^{ij}$. We say that the set M is* satisfiable *if there exists an $X \neq \emptyset$ such that M is satisfiable modulo X.*

Definition 5 (Mapping entailment). *Let O_i and O_j be two ontologies and let X be a non-empty finite set of fresh individual constants. Also, let M be a set of weighted mappings from O_i to O_j. The set M* entails $i{:}C\ r_{[a,b]}\ j{:}D$ modulo X*, denoted $M \models_X i{:}C\ r_{[a,b]}\ j{:}D$, if for every interpretations $\mathcal{I}_i$ and $\mathcal{I}_j$ of O_i and O_j, respectively, such that $(\mathcal{I}_i, \mathcal{I}_j)$ satisfies M modulo X, we have that $(\mathcal{I}_i, \mathcal{I}_j) \models_X i{:}C\ r_{[a,b]}\ j{:}D$. The set M* entails $i{:}C\ r_{[a,b]}\ j{:}D$*, in symbols, $M \models i{:}C\ r_{[a,b]}\ j{:}D$, if we have $M \models_X i{:}C\ r_{[a,b]}\ j{:}D$ for every $X \neq \emptyset$.*

Remark 2 (Inconsistent mappings). In the case of $a > b$, the mapping $i{:}C\ r_{[a,b]}\ j{:}D$ has no satisfying interpretations. That is, mappings defined over empty probability ranges are inconsistent mappings and we denote them by FALSE. This allows us to express unsatisfiability of a set of mappings in terms of mapping entailment: if M is a set of mappings between O_i to O_j, then $M \models FALSE$ is equivalent to stating that M is not satisfiable. Moreover, notice that $FALSE \models i{:}C\ r_{[a,b]}\ j{:}D$, i.e. every mapping is entailed by the inconsistent mapping.

3 Adequacy and Expressivity of the Semantics

In this section we prove the adequacy of the formal semantics for weighted mappings proposed in Section 2 by showing that (i) it is a conservative extension of a standard semantics for crisp mappings; (ii) it can be used to provide a formal interpretation of the results returned by automatic ontology matching algorithms; and (iii) it is general enough to provide a uniform formal interpretation of weighted mappings between both concepts and individuals.

3.1 Backward Compatibility with the Semantics of Crisp Mappings

In order to prove that the classificational semantics presented in this paper provides a suitable extension of the one for crisp mappings we show that it is a conservative extension of the DDL-based semantics for crisp mappings presented in [3] when the intervals $[a, b]$ are either $[0, 0]$ or $[1, 1]$. First of all, notice that when a and b are either 0 or 1 every mapping can be expressed in terms of $\sqsubseteq_{[1,1]}$. Indeed the following logical consequences hold:

- $i{:}C \sqsupseteq_{[1,1]} j{:}D$ is equivalent to $j{:}D \sqsubseteq_{[1,1]} i{:}C$
- $i{:}C \equiv_{[1,1]} j{:}D$ is equivalent to $i{:}C \sqsubseteq_{[1,1]} j{:}D$ and $j{:}D \sqsubseteq_{[1,1]} i{:}C$
- $i{:}C \perp_{[1,1]} j{:}D$ is equivalent to $i{:}C \sqsubseteq_{[1,1]} j{:}\neg D$
- $i{:}C \sqsubseteq_{[0,0]} j{:}D$ is equivalent to $i{:}C \sqsubseteq_{[1,1]} j{:}\neg D$
- $i{:}C \sqsupseteq_{[0,0]} j{:}D$ is equivalent to $i{:}C \sqsubseteq_{[1,1]} j{:}\neg D$
- $i{:}C \equiv_{[0,0]} j{:}D$ is equivalent to $i{:}C \sqsubseteq_{[1,1]} j{:}\neg D$

- $i{:}C \perp_{[0,0]} j{:}D$ is equivalent to $i{:}C \sqsubseteq_{[1,1]} j{:}D$ and $j{:}D \sqsubseteq_{[1,1]} i{:}C$
- $i{:}C\ r_{[0,1]}\ j{:}D$ is equivalent to $i{:}\bot \sqsubseteq_{[1,1]} j{:}\top$
- $i{:}C\ r_{[1,0]}\ j{:}D$ is equivalent to $i{:}\top \sqsubseteq_{[1,1]} j{:}\bot$

The above guarantees that any weighted mapping from O_i to O_j with weight in $\{[0,0],[0,1],[1,0],[1,1]\}$ can be rewritten as a $\sqsubseteq_{[1,1]}$-mapping between O_i and O_j (i.e. from O_i to O_j or from O_j to O_i). We say that a set of weighted mappings between O_i and O_j is in $\sqsubseteq_{[1,1]}$-normal form if they are of the form $x{:}C \sqsubseteq_{[1,1]} y{:}D$.

Lemma 1. *If M is a set of mappings between two ontologies O_i and O_j with weights in $\{[0,0],[0,1],[1,0],[1,1]\}$, then M can be rewritten in an equivalent set of mappings $M_{\sqsubseteq_{[1,1]}}$ in $\sqsubseteq_{[1,1]}$-normal form.*

Crisp mappings are defined in DDL via bridge rules, whose syntax and semantics are as follows: let $\mathcal{I}_i$ and $\mathcal{I}_j$ be interpretations of the ontologies O_i and O_j, resp. Let $\rho_{ij} \subseteq \Delta^{\mathcal{I}_i} \times \Delta^{\mathcal{I}_j}$ be a domain correspondence relation. In DDL we define four kinds of bridge rules, but, due to the rewriting described above, the only one that is interesting here is what the so-called into bridge rule:

$$(\mathcal{I}_i, \mathcal{I}_j, \rho_{ij}) \models i{:}C \xrightarrow{\sqsubseteq} j{:}D \quad \text{iff} \quad \rho_{ij}(C^{\mathcal{I}_i}) \subseteq D^{\mathcal{I}_j}$$

Lemma 2. *Let O_i and O_j be two ontologies.*
(1) Let $\mathcal{I}_i$ and $\mathcal{I}_j$ be two interpretations of O_i and O_j, resp., extended with a non-empty finite X. Then there exists a domain relation $\rho_{ij} \subseteq \Delta^{\mathcal{I}_i} \times \Delta^{\mathcal{I}_j}$ such that

$$(\mathcal{I}_i, \mathcal{I}_j) \models_X i{:}C \sqsubseteq_{[1,1]} j{:}D \quad \textit{iff} \quad (\mathcal{I}_i, \mathcal{I}_j, \rho_{ij}) \models i{:}C \xrightarrow{\sqsubseteq} j{:}D$$

(2) Let $\mathcal{I}_i$ and $\mathcal{I}_j$ be interpretations of O_i and O_j with finite domains *$\Delta^{\mathcal{I}_i}$ and $\Delta^{\mathcal{I}_j}$, and let $\rho_{ij} \subseteq \Delta^{\mathcal{I}_i} \times \Delta^{\mathcal{I}_j}$ be a* DDL *domain relation. Then there exists a non-empty finite set X of fresh individuals and interpretations $\mathcal{I}'_i$ and $\mathcal{I}'_j$ extending $\mathcal{I}_i$ and $\mathcal{I}_j$, resp., over X such that*

$$(\mathcal{I}_i, \mathcal{I}_j, \rho_{ij}) \models i{:}C \xrightarrow{\sqsubseteq} j{:}D \quad \textit{iff} \quad (\mathcal{I}'_i, \mathcal{I}'_j) \models_X i{:}C \sqsubseteq_{[1,1]} j{:}D$$

Proof (Outline). For (1) we define $\rho_{ij} = \{\langle x^{\mathcal{I}_i}, x^{\mathcal{I}_j}\rangle \mid x \in X\}$ and then show that $R(C_X^{\mathcal{I}_i}, D_X^{\mathcal{I}_j}) = 1$ iff $\rho_{ij}(C^{\mathcal{I}_i}) \subseteq D^{\mathcal{I}_j}$. In order to prove (2) we first choose a set of fresh constants $X = \{x_{(c,d)} \mid (c,d) \in \rho_{ij}\}$ (i.e one constant for each pair of the domain relation), and then extend $\mathcal{I}_i$ and $\mathcal{I}_j$ to $\mathcal{I}'_i$ and $\mathcal{I}'_j$ over X by defining $\mathcal{I}'_i(x_{(c,d)}) = c$ and $\mathcal{I}'_j(x_{(c,d)}) = d$. Then, $\rho_{ij}(C^{\mathcal{I}_i}) \subseteq D^{\mathcal{I}_j}$ iff $R(C_X^{\mathcal{I}'_i}, D_X^{\mathcal{I}'_j}) = 1$.

Theorem 1. *Consider ontologies O_i and O_j in ontology languages that have the finite-model property (recall that $\mathcal{ALCO}$ has this property). Let M be a set of mappings with weights in $\{[0,0],[0,1],[1,0],[1,1]\}$ and $M_{\sqsubseteq_{[1,1]}}$ its $\sqsubseteq_{[1,1]}$-normal form. Let $M_{\text{DDL}} = \{x{:}C \xrightarrow{\sqsubseteq} y{:}D \mid x{:}C \sqsubseteq_{[1,1]} y{:}D \in M_{\sqsubseteq_{[1,1]}}\}$ be the set of* corresponding DDL *bridge rules. Then, for arbitrary concepts C and D*

$$M \models i{:}C \sqsubseteq_{[1,1]} j{:}D \quad \textit{iff} \quad M_{\text{DDL}} \models_{inv} i{:}C \xrightarrow{\sqsubseteq} j{:}D$$

where $\models_{inv}$ denotes DDL *logical consequence restricted to models where $\rho_{ij} = \rho_{ji}^{-1}$.*

Proof (outline). Let us assume that $M_{\sqsubseteq_{[1,1]}} \not\models i{:}C \sqsubseteq_{[1,1]} j{:}D$. Then there exists a set $X \neq \emptyset$ and interpretations $\mathcal{I}_i$ and $\mathcal{I}_j$ of O_i and O_j, resp., extended with X such that $(\mathcal{I}_i, \mathcal{I}_j) \models_X M_{\sqsubseteq_{[1,1]}}$ but $(\mathcal{I}_i, \mathcal{I}_j) \not\models_X i{:}C \sqsubseteq_{[1,1]} j{:}D$. By Lemma 2, there is a domain relation ρ_{ij} for which $(\mathcal{I}_i, \mathcal{I}_j, \{\rho_{ij}, \rho_{ji}\}) \models M_{\text{DDL}}$ and $(\mathcal{I}_i, \mathcal{I}_j, \rho_{ij}) \not\models i{:}C \xrightarrow{r} j{:}D$. Therefore, $M_{\text{DDL}} \not\models i{:}C \xrightarrow{\sqsubseteq} j{:}D$.

Vice versa, assume $M_{\text{DDL}} \not\models i{:}C \xrightarrow{\sqsubseteq} j{:}D$. Thus, there is a DDL interpretation $(\mathcal{I}_i, \mathcal{I}_j, \{\rho_{ij}, \rho_{ji}\})$ with $\rho_{ij} = \rho_{ji}^{-1}$ such that $(\mathcal{I}_i, \mathcal{I}_j, \{\rho_{ij}, \rho_{ij}\}) \models M_{\text{DDL}}$ whereas $(\mathcal{I}_i, \mathcal{I}_j, \rho_{ij}) \not\models i{:}C \xrightarrow{\sqsubseteq} j{:}D$. One can extend the finite domain property of most DLs supporting qualified existential restrictions to those of DDLs using them via the "global DL" construction given in [3]. Therefore, we can suppose w.l.o.g that the domains of $\mathcal{I}_i$ and $\mathcal{I}_j$ are finite and, thus, ρ_{ij} and ρ_{ij} are also finite. By lemma 2, we have that there is an X such that $(\mathcal{I}_i, \mathcal{I}_j) \models_X M_{\sqsubseteq_{[1,1]}}$ but $(\mathcal{I}_i, \mathcal{I}_j) \not\models_X i{:}C \sqsubseteq_{[1,1]} j{:}D$. Notice that the fact that $\rho_{ij} = \rho_{ji}^{-1}$ guarantees that the Xs associated to ρ_{ij} and ρ_{ji} are the same.

3.2 Interpreting the Results of Ontology Matchers

The semantics presented in Section 2 expresses the weight of a mapping between two elements by means of an interval $[a, b]$, while ontology matching algorithms usually return a single confidence value c. Thus, we need to ask ourselves how we can represent this value c by means of the interval $[a, b]$. In answering this question we can opt for several alternatives: we can decide to represent c by means of the (pointwise) interval $[c, c]$, or the interval $[0, c]$, or the interval $[c, 1]$, or the interval $[c - \epsilon, c + \epsilon]$ centered in c. This flexibility of representation allows us to capture the different assumptions which are used by the different algorithms. Let us illustrate this by means of some examples. If we take a low confidence value c, some algorithms interpret it as an "I don't know" answer; others, instead, use it to represent the fact that "the two concepts are very different". These two usages of c can be captured in our formalism by two different encodings: in the first case c is formalized by the interval $[c, 1]$; in the second c corresponds to the interval $[0, c]$. A possible different representation is given when the result of an ontology matcher expresses an estimation of similarity with some degree of approximation. In this case the returned value c can be represented by means of the centered interval $[\max(0, c - \epsilon), \min(c + \epsilon, 1)]$, where ϵ is a value between 0 and 0.5 that depends on the level of accuracy of the matching algorithm: the more accurate the matcher is, the smaller the ϵ will be. As we can see from these few examples, the representation of c by means of the pointwise interval $[c, c]$ is only one among a set of different choices, and a very challenging one, since it says that the ontology matching algorithm returns the exact level of matching between two elements with a perfect level of accuracy (i.e. $\epsilon = 0$).

3.3 Uniform Semantics for Mapping between Concepts and Individuals

The semantics presented in Section 2 gives a uniform framework to interpret mappings that involve pairs of concepts, and pairs of individuals. If $C_X^{\mathcal{I}}$ is the set of elements of X that could be reclassified in C, $a_X^{\mathcal{I}}$ is the set of elements of X that could be the same

as a. Contrary to $a^{\mathcal{I}}$ which is an element of the domain $\Delta_{\mathcal{I}}$, $a^{\mathcal{I}}_X$ is a subset of X. For this reason, we use the same notation for individual as for classes: $a^{\mathcal{I}}_X = \{a\}^{\mathcal{I}}_X$ and we use the notation $1:\{a\} \equiv_{[0.8,1]} 2:\{b\}$ to represent the fact a is "almost the same as" b. According to the formal semantics we have that

$$(\mathcal{I}_1, \mathcal{I}_2) \models_X 1:\{a\} \equiv_{[0.8,1]} 2:\{b\} \quad \text{iff} \quad 0.8 \leq 2 * \frac{|\{a\}^{\mathcal{I}_1}_X \cap \{b\}^{\mathcal{I}_2}_X|}{|\{a\}^{\mathcal{I}_1}_X| + |\{b\}^{\mathcal{I}_2}_X|} \leq 1$$

Therefore, the mapping $\{a\} \equiv_{[0.8,1]} \{b\}$ states that the harmonic mean of the fraction of items in X equivalent to a that are also equivalent to b, and the fraction of items in X equivalent to b which are also equivalent to a is between the numbers 0.8 and 1.

In order to better understand the intuition behind this formalization, consider the individuals Trento and TrentoTown belonging to two ontologies O_1 and O_2, resp. While there are interpretations $\mathcal{I}_1$ and $\mathcal{I}_2$ in which the two individuals coincide, that is, $\{\mathsf{Trento}\}^{\mathcal{I}_1}_X = \{\mathsf{TrentoTown}\}^{\mathcal{I}_2}_X$, there may be cases in which Trento is considered to be an area broader than TrentoTown but still largely overlapping with it. This can be formalized by considering two interpretations $\mathcal{I}_1$ and $\mathcal{I}_2$ where $\{\mathsf{Trento}\}^{\mathcal{I}_1}_X$ includes $\{\mathsf{TrentoTown}\}^{\mathcal{I}_2}_X$, for instance, $\{\mathsf{Trento}\}^{\mathcal{I}_1}_X = \{x_1, x_2, x_3, x_4\}$ and $\{\mathsf{TrentoTown}\}^{\mathcal{I}_2}_X = \{x_1, x_2, x_3\}$. In this case the mapping is weighted as

$$(\mathcal{I}_1, \mathcal{I}_2) \models_X 1:\{\mathsf{Trento}\} \equiv_{[0.85,0.86]} 2:\{\mathsf{TrentoTown}\}$$

where values 0.85 and 0.86 are obtained as under and over approximation of the fraction

$$2 * \frac{\left|\{\mathsf{Trento}\}^{\mathcal{I}_1}_X \cap \{\mathsf{TrentoTown}\}^{\mathcal{I}_2}_X\right|}{\left|\{\mathsf{Trento}\}^{\mathcal{I}_1}_X\right| + \left|\{\mathsf{TrentoTown}\}^{\mathcal{I}_2}_X\right|} = 2 * \frac{3}{7} = \frac{6}{7}$$

4 Properties of Mapping Entailment

In this section, we show how the proposed semantics can be used to compute additional mappings which are logical consequences of an initial set of mappings. The proofs are omitted for lack of space and can be found in [1].

Proposition 1 shows general properties of mapping entailment independently of the mapping relation considered.

Proposition 1. *The following hold:*

1. $\models i:A\ r_{[0,1]}\ j:G$
2. $\mathsf{FALSE} \models i:A\ r_{[a,b]}\ j:G$
3. $i:A\ r_{[a,b]}\ j:G \models i:A\ r_{[c,d]}\ j:G$ *if* $[a,b] \subseteq [c,d]$.
4. $i:A\ r_{[a,b]}\ j:G, i:A\ r_{[c,d]}\ j:G \models i:A\ r_{[v,w]}\ j:G$, $v = \max(a,c)$, $w = \min(b,d)$.
5. $i:A\ r_{[a,b]}\ j:G \models j:G\ r^{-1}_{[a,b]}\ i:A$, *where* r^{-1} *is the inverse relation of* r.[6]

[6] Recall that $\equiv^{-1}$ is $\equiv$, $\sqsubseteq^{-1}$ is $\sqsupseteq$, $\sqsupseteq^{-1}$ is $\sqsubseteq$, and $\bot^{-1}$ is $\bot$.

Proposition 2 includes properties of mapping entailment which relate equivalence with subsumption and disjointness.

Proposition 2. *The following hold:*

1. $i\!:\!A \sqsubseteq_{[a,b]} j\!:\!G,\ i\!:\!A \sqsupseteq_{[c,d]} j\!:\!G \models i\!:\!A \equiv_{[v,w]} j\!:\!G$ *where*

$$v = \begin{cases} \frac{2ac}{a+c} & \text{if } a \neq 0 \text{ or } c \neq 0 \\ 0 & \text{if } a = c = 0 \end{cases} \quad \textit{and} \quad w = \begin{cases} \frac{2bd}{b+d} & \text{if } b \neq 0 \text{ or } d \neq 0 \\ 0 & \text{if } b = d = 0 \end{cases}$$

2. $i\!:\!A \equiv_{[a,b]} j\!:\!G \models i\!:\!A \sqsubseteq_{[v,1]} j\!:\!G$ *and*
   $i\!:\!A \equiv_{[a,b]} j\!:\!G \models i\!:\!A \sqsupseteq_{[v,1]} j\!:\!G$ *where* $v = \frac{a}{2-a}$
3. $i\!:\!A \equiv_{[a,b]} j\!:\!G,\ i\!:\!A \sqsubseteq_{[c,d]} j\!:\!G \models i\!:\!A \sqsupseteq_{[v,w]} j\!:\!G$ *and*
   $i\!:\!A \equiv_{[a,b]} j\!:\!G,\ i\!:\!A \sqsupseteq_{[c,d]} j\!:\!G \models i\!:\!A \sqsubseteq_{[v,w]} j\!:\!G$ *where*

$$v = \begin{cases} \frac{ac}{2d-a} & \text{if } a \neq 2d \\ 0 & \text{otherwise} \end{cases} \quad \textit{and} \quad w = \begin{cases} 0 & \text{if } b = 0 \\ \frac{bd}{2c-b} & \text{if } b \neq 0 \text{ and } b < \frac{2c}{1+d} \\ 1 & \text{otherwise} \end{cases}$$

4. $i\!:\!A \equiv_{[a,b]} j\!:\!G \models i\!:\!A \perp_{[v,w]} j\!:\!G$ *and*
   $i\!:\!A \perp_{[a,b]} j\!:\!G \models i\!:\!A \equiv_{[v,w]} j\!:\!G$ *where* $v = 1 - b$ *and* $w = 1 - a$

Proposition 3 introduces properties of mapping entailment with respect to $\mathcal{ALCO}$ constructors. No property about existential restriction is included, since, as mentioned before, mappings between roles are not considered in this paper.

Proposition 3. *The following hold:*

1. $i\!:\!A \sqsubseteq_{[a,b]} j\!:\!G \models i\!:\!A \sqsubseteq_{[v,w]} j\!:\!\neg G$ *where* $v = 1 - b$ *and* $w = 1 - a$
2. $\left.\begin{array}{l} i\!:\!A \sqsubseteq_{[a,b]} j\!:\!G \\ i\!:\!A \sqsupseteq_{[c,d]} j\!:\!\top \\ i\!:\!\top \sqsubseteq_{[e,f]} j\!:\!G \end{array}\right\} \models i\!:\!\neg A \sqsubseteq_{[v,w]} j\!:\!G$ *where*

$$v = \begin{cases} \max\left(0, \frac{e-bd}{1-c}\right) & \text{if } c \neq 1 \\ 1 & \text{if } c = 1 \end{cases} \quad \textit{and} \quad w = \begin{cases} \min\left(\frac{f-ac}{1-d}, 1\right) & \text{if } d \neq 1 \\ 1 & \text{if } d = 1 \end{cases}$$

3. $\left.\begin{array}{l} i\!:\!A \sqsubseteq_{[a,b]} j\!:\!G \\ i\!:\!A \sqsubseteq_{[c,d]} j\!:\!H \end{array}\right\} \models i\!:\!A \sqsubseteq_{[v,w]} j\!:\!G \sqcap H$ *where* $\begin{cases} v = \max(0, a + c - 1), \textit{ and} \\ w = \min(b, d) \end{cases}$
4. $\left.\begin{array}{l} i\!:\!A \sqsubseteq_{[a,b]} j\!:\!G \\ i\!:\!A \sqcap B \sqsupseteq_{[c,d]} j\!:\!G \end{array}\right\} \models i\!:\!A \sqcap B \sqsubseteq_{[v,1]} j\!:\!G$ *where* $v = ac$
5. $i\!:\!A \sqsubseteq_{[a,b]} j\!:\!G \sqcap H \models i\!:\!A \sqsubseteq_{[a,1]} j\!:\!G$
6. $\left.\begin{array}{l} i\!:\!A \sqcap B \sqsubseteq_{[a,b]} j\!:\!G \\ i\!:\!A \sqcap B \sqsupseteq_{[c,d]} j\!:\!G \end{array}\right\} \models i\!:\!A \sqsubseteq_{[0,w]} j\!:\!G$ *where* $w = \frac{b}{c}$
7. $i\!:\!A \ \sqsubseteq_{[a,b]}\ j\!:\!\{g\} \models i\!:\!A \ \sqsubseteq_{[0,b]}\ j\!:\!\{g\} \sqcap G$
8. $i\!:\!\{a\} \ \sqsubseteq_{[a,b]}\ j\!:\!G \models i\!:\!\{a\} \sqcap A \ \sqsubseteq_{[a,1]}\ j\!:\!G$

Notice that proposition 3 does not contain any rule that allow to infer some weight interval $[v, w]$ of the mapping $1 : A \sqcap B \ \sqsubseteq_{[v,w]} \ 2 : C$ starting from the weight intervals $[a, a']$ and $[b, b']$ of the mappings $i : A \ \sqsubseteq_{[a,a']} \ 2 : C$ and $i : B \ \sqsubseteq_{[b,b']} \ 2 : C$. This contrasts to what happens for crisp mappings, where $1 : A \sqcap B \ \sqsubseteq_{[1,1]} \ 2 : C$ is a logical consequence of the two mappings $i : A \ \sqsubseteq_{[1,1]} \ 2 : C$ and $i : B \ \sqsubseteq_{[1,1]} \ 2 : C$. In general there is an independence between the weights associated to the mapping of two concepts and the weight associated to the mapping of their conjunction. Example 3 provides an evidence of this independence by showing that mappings of the form $i : A \ \sqsubseteq_{[x,x]} \ 2 : C$ and $i : B \ \sqsubseteq_{[y,y]} \ 2 : C$ with high (resp. low) values of x and y are consistent with mappings of the form $1 : A \sqcap B \ \sqsubseteq_{[z,z]} \ 2 : C$ with a low (resp. high) value for z.

Example 3. Suppose that O_1 contains the concepts Professor and Professional, while O_2 includes the concepts HasFreeTime and EarnsALot. The following two interpretations $\mathcal{I}_1$ and $\mathcal{I}_2$

$$\mathsf{Professor}^{\mathcal{I}_1}_X = \{x_1, \ldots, x_{100}\} \qquad \mathsf{HasFreeTime}^{\mathcal{I}_2}_X = \{x_1, \ldots, x_{90}, x_{101}, \ldots, x_{190}\}$$
$$\mathsf{Professional}^{\mathcal{I}_1}_X = \{x_{91}, \ldots, x_{190}\} \quad \mathsf{EarnsALot}^{\mathcal{I}_2}_X = \{x_{91}, \ldots, x_{100}\}$$

satisfies the following mappings:

$$(\mathcal{I}_1, \mathcal{I}_2) \models_X 1{:}\mathsf{Professor} \sqsubseteq_{[0.9,0.9]} 2{:}\mathsf{HasFreeTime} \tag{3}$$
$$(\mathcal{I}_1, \mathcal{I}_2) \models_X 1{:}\mathsf{Professional} \sqsubseteq_{[0.9,0.9]} 2{:}\mathsf{HasFreeTime} \tag{4}$$
$$(\mathcal{I}_1, \mathcal{I}_2) \models_X 1{:}\mathsf{Professor} \sqcap \mathsf{Professional} \sqsubseteq_{[0.1,0.1]} 2{:}\mathsf{HasFreeTime} \tag{5}$$
$$(\mathcal{I}_1, \mathcal{I}_2) \models_X 1{:}\mathsf{Professor} \sqsubseteq_{[0.1,0.1]} 2{:}\mathsf{EarnsALot} \tag{6}$$
$$(\mathcal{I}_1, \mathcal{I}_2) \models_X 1{:}\mathsf{Professional} \sqsubseteq_{[0.1,0.1]} 2{:}\mathsf{EarnsALot} \tag{7}$$
$$(\mathcal{I}_1, \mathcal{I}_2) \models_X 1{:}\mathsf{Professor} \sqcap \mathsf{Professional} \sqsubseteq_{[0.9,0.9]} 2{:}\mathsf{EarnsALot} \tag{8}$$

Notice that $\mathcal{I}_1$ and $\mathcal{I}_2$ satisfy the low weight mapping (5), on the conjunction of two concepts, as well as the high weight mappings (3) and (4) defined on the two conjunct concepts. Conversely, the two interpretations satisfy the high weight mapping (8) and the low weight mapping (6) and (7).

Proposition 4 shows mapping entailments in the presence of local knowledge. Local entailment in ontology O_i is denoted by $\models_i$. As in the case of weighted mappings, if we write $\models_i C \ r \ D$, we assume that C and D belong to $L(O_i)$.

Proposition 4. *The following hold:*

1. *If* $\models_i A \equiv B$ *and* $\models_j G \equiv H$ *then* $i{:}A \ r_{[v,w]} \ j{:}G \models i{:}B \ r_{[v,w]} \ j{:}H$
2. *If* $\models_i A \sqsubseteq B$ *then* $i{:}A \sqsupseteq_{[v,w]} j{:}G \models i{:}B \sqsupseteq_{[v,1]} j{:}G$
3. *If* $\models_i A \sqsupseteq B$ *then* $i{:}A \sqsupseteq_{[v,w]} j{:}G \models i{:}B \sqsupseteq_{[0,w]} j{:}G$
4. *If* $\models_i A \sqcap B \sqsubseteq \bot$ *then* $i{:}A \sqsupseteq_{[v,w]} j{:}G \models i{:}B \sqsupseteq_{[0,1-v]} j{:}G$

5 Variations on a Theme

There are a number of places above where we have made certain choices that could have been done differently. We examine some of the alternatives in this section.

5.1 Interpreting Weighted Equivalence Mappings

The choice of interpreting the weighted equivalence mapping by means of the F-measure is based on the fact that this is the typical measure used to evaluate the global quality of a classifier. However, any function $f : [0,1]^2 \rightarrow [0,1]$ that satisfies the following properties can be chosen to combine specific precision and recall values:

1. $f(0,0) = f(1,0) = f(0,1) = 0$, and $f(1,1) = 1$. The motivation for this is so that in the case of crisp mappings, the weight of an equivalence mapping be the logical "and" of the weights of the subsumption mappings.
2. $f(\cdot,\cdot)$ is monotonic in each variable. In other words, if the weight of a subsumption mapping increases then the value of the equivalence mapping must also increase.

In this paper, $f(\cdot,\cdot)$ is the F-measure, but parallel studies can be conducted where, for instance, $f(x,y) = \max(x,y)$, $f(x,y) = x \cdot y$ or $f(x,y) = \frac{x+y}{2}$.

A different approach to interpret the weighted equivalence mapping, starts from the usual definition of equivalence in DLs in terms of subsumption: ($A \equiv B$) iff ($A \sqsubseteq B$) and ($A \sqsupseteq B$). When dealing with single numbers for precision and recall, it is usually imposible to combine them into a single value by simple conjunction; hence the use of F as above. However, when using *ranges* of scores $[a,b]$ for subsumption relations, one can define $A \equiv^{\vee}_{[a,b]} B$ as the conjunction of $A \sqsubseteq_{[a,b]} B$ and $A \sqsupseteq_{[a,b]} B$. This leads to a natural rule: if $A \sqsubseteq_{[a_1,b_1]} B$ and $A \sqsupseteq_{[a_2,b_2]} B$ then $A \equiv^{\vee}_{[v,w]} B$ with $v = \min(a_1,a_2)$ and $w = \max(b_1,b_2)$.

The current version of our semantics is characterised by the fact that it allows a single individual in ontology O_1 to be matched with certainty to *sets* of individuals in O_2. That is, we can express $\{1{:}book1\} \equiv_{[1,1]} \{2{:}book1_copy1, 2{:}book1_copy2\}$. This is certainly useful if we want to express the fact that a book (e.g. $book1$) is equivalent to all the copies of that book. On the other hand, as pointed out by Cuenca Grau et al [7] the DDL which corresponds to our semantics, constructed according to Theorem 1 has certain conclusions that may sometimes be undesirable. The example in [7] starts from an ontology O_1 which states that the concepts Flying and NoNFlying are disjoint and that all birds can fly, and an ontology O_2 which defines the concept Penguin, and it connects the two ontologies with DDL bridge rules which map Bird onto Penguin and NoNFlying onto Penguin as described below:

$$1{:}\mathsf{NoNFlying} \sqcap \mathsf{Flying} \sqsubseteq \bot, \qquad 1{:}\mathsf{Bird} \xrightarrow{\sqsupseteq} 2{:}\mathsf{Penguin},$$

$$1{:}\mathsf{Bird} \sqsubseteq \mathsf{Flying}, \qquad 1{:}\mathsf{NoNFlying} \xrightarrow{\sqsupseteq} 2{:}\mathsf{Penguin}$$

$$2{:}\mathsf{Penguin} \sqsubseteq \top,$$

The point made in [7] is that if we describe this example in one ontology, by rewriting $\xrightarrow{\sqsubseteq}$ and $\xrightarrow{\sqsupseteq}$ mappings by means of $\sqsubseteq$ and $\sqsupseteq$ statements, we obtain an unsatisfiable Penguin concept, as penguins cannot be (flying) birds and non flying creatures at the same time. In DDL the two ontologies O_1, O_2 and the above mappings are satisfiable. Intuitively this happens because a penguin x in ontology O_2 can correspond, via the DDL domain relation, to two distinct objects in O_1: one flying, and the other not.

This effect can be avoided by restricting the domain relation in DDL to be 1-to-1, as in P-DL [2]. A similar restriction can also be imposed to the semantics of weighted mappings given in Section 2: one simply requires that the interpretations $\mathcal{I}_1$ and $\mathcal{I}_2$ be 1-to-1 on the set X.[7] It can easily be verified that the corresponding proofs of Lemma 2 and of Theorem 1 go through. Note that in the case of 1-to-1 relations, all satisfiable mappings between individuals can be reduced to mappings in the intervals $[0, 0]$, $[1, 1]$ or $[0, 1]$. That because the value of precision, recall, and F-measure in this case is either 0 or 1.

5.2 A General Framework for Probabilistic Mappings

A natural question to ask is why not simply union the two ontologies O_1 and O_2 with the precision and recall statements expressed as probabilistic subsumptions between their terms, into one "global" probabilistic DL ontology (PDLO), and then reason with it. Our answer is that intuitively one wants to keep their domains of interpretation disjoint. The set X, whose identifiers did not appear in either O_1 nor O_2, but were independently interpreted into the domains of O_1 and O_2 respectively, played a crucial role in this.

We show here how one could create a global PDLO which respects this intuition, and from which one could draw a variety of conclusions depending on which specific probabilistic DL was chosen. (See Lukasiewicz and Straccia's review [13] for a variety of proposals.) The idea of such a translation is inspired by our earlier work on DDL [3], where we also constructed a single global DL, though the details of our construction here are different.

Let us suppose that O_1 and O_2 are two ontologies and let $P = \{1\!:\!A_k \sqsubseteq_{[a,b]} 2\!:\!G_k\}$ be a set of weighted inclusion mapping statements between O_1 and O_2.

Suppose that $\mathcal{DL}_\mathcal{P}$ is some probabilistic description logic, where one can make inequality assertions on the probability of subsumptions of the form (E $\sqsubseteq$ F) *with probability* $\lesseqgtr p$.

Then consider the following translation from O_1, O_2 and P into a $\mathcal{DL}_\mathcal{P}$ TBox T_{12}:

- The atomic concept symbols of T_{12} consist of $\{1\!:\!A \mid A \in CN_1\} \cup \{2\!:\!G \mid G \in CN_2\} \cup \{1\!:\!\mathsf{ANYTHING}, 2\!:\!\mathsf{ANYTHING}, X\}$.
- The atomic role symbols of T_{12} consist of $\{1 : R \mid R \in RN_1\} \cup \{2 : S \mid S \in RN_2\} \cup \{\rho_1, \rho_2\}$.
- The axioms of T_{12} include
  - relabelled local axioms of O_1 and O_2, obtained by prefixing all identifiers and subconcepts with 1 and 2 respectively.
  - transformations of all concepts and axioms, as described in [3], so that complementation in O_i is with respect to $i\!:\!\mathsf{ANYTHING}$ only (e.g., $\neg 1\!:\!A$ is replaced by $1\!:\!\mathsf{ANYTHING} \sqcap \neg 1\!:\!A$).
  - $(1 : \mathsf{ANYTHING} \sqcap 2 : \mathsf{ANYTHING} \sqsubseteq \bot)$, expressing the disjointness of the domains of the two ontologies being mapped.
  - axioms restricting ρ_i to have domain X and range $i\!:\!\mathsf{ANYTHING}$ for $i = 1, 2$. The idea is that ρ_i will play the role of the interpretation $\mathcal{I}_i$ applied to X.

[7] This corresponds to requiring the values in X to obey the unique name assumption.

- for every statement $A \sqsubseteq_{[a,b]} G$ add $(\exists\rho_1.1 : A \sqsubseteq \exists\rho_2.2 : G)$ *with probability* $\geq a$ and $(\exists\rho_1.1 : A \sqsubseteq \exists\rho_2.2 : G)$ *with probability* $\leq b$
  Essentially, these axioms establish the probabilistic subset relationships between the elements of X mapped by ρ_i
- converse axioms for every statement $A \sqsupseteq_{[a,b]} G$, adding $(\exists\rho_2.2 : G \sqsubseteq \exists\rho_1.1 : A)$ *with probability* $\geq a$ and $(\exists\rho_2.2 : G \sqsubseteq \exists\rho_1.1 : A)$ *with probability* $\leq b$
- If $\mathcal{DL}_\mathcal{P}$ requires all axioms to be stated with a probability, add probability 1 to the axioms built before the last 2 steps.
- In case one wants to simulate the approach presented in this paper, one would also like to enforce that X is a finite set. If the $\mathcal{DL}_\mathcal{P}$ has the finite model property, then we are safe since we can restrict ourselves to reasoning in finite models. Otherwise, if X could be made finite using axioms then so could the domain concept, $\top$, and it is known that reasoning in only finite models leads to very different and complex deductions [4].

We believe the above framework (obviously adjusted for differences like the use of Bayes nets for probabilistic statements) will allow us to explore in the future the differences/similarities and benefits/tradeoffs between various ways of defining weighted mappings based on the numerous proposals for probabilistic DLs in the literature, as well as the one introduced in this paper. The first candidate for such an investigation will be Heinsohn's pioneering proposal for probabilistic DL, $\mathcal{ALCP}$ [9]. Heinsohn models uncertain subsumption between concepts as conditional probabilities — what he calls p-conditionings — and also considers intervals instead of single values.

6 Related Work

In the last years there has been a growing interest in the management of uncertainty and vagueness in DLs for the semantic web [13]. Two main lines of research have been followed: probabilistic generalisations of DLs to deal with uncertainty, and also fuzzy extensions of DLs to handle vagueness. Since our approach falls into the first category we only report, in the presence of comprehensive surveys such that [13], some of the previous attempts to combine DLs with probability, paying special attention to those related to ontology mapping.

Heinsohn was one of the first to provide a probabilistic extension of description logics [9]. The language $\mathcal{ACLP}$ builds over $\mathcal{ACL}$ and adds probabilistic subsumptions between concepts. These are formalised in terms of the so-called p-conditionings which encode that the conditional probability of a concept given another concept lies in a concrete real interval. Heinsohn assumes interpretation domains to be finite, and defines conditional probabilities in terms of set cardinalities.

Koller et al. presented P-CLASSIC [10] which is a probabilistic generalisation of a version of the description logic CLASSIC. P-CLASSIC follows Heinsohn's approach and aims at answering probabilistic subsumption queries, but its semantics is based on a reduction to Bayesian networks.

More recently, Lukasiewicz proposed probabilistic extensions of $\mathcal{SHOIN}(\mathbf{D})$ and $\mathcal{SHIF}(\mathbf{D})$ [11]. Uncertain knowledge is realised by way of conditional constraints

which, similarly to p-conditionings, encode interval restrictions for conditional probabilities over concepts. The semantics is based on the notion of lexicographic entailment in probabilistic default reasoning. Conditional constraints can be applied to individuals too. Thus, it is possible to represent, besides statistical knowledge about concepts, degrees of belief about individuals.

Although these approaches do not tackle the formalisation of ontology mappings directly, they certainly could be used for this purpose. Given two ontologies O_1 and O_2 one could express probabilistic subsumptions between their concepts within one "global" probabilistic ontology (in one of the formalisms described above), and then reason with it. The main motivation of this work is to provide a formalism that keeps the domains of the ontology interpretations disjoint, as it is done in [3, 16] for crisp mappings. The classificational interpretation of mappings is itself a novel approach and justifies the use of F-measure to model equivalence mappings which could not be addressed with most of the existing probabilistic description logics.

From a very different perspective, Lukasiewicz et al. presents in [12] a language for representing and reasoning with uncertain ontology mappings. This approach is based on the tight integration of disjunctive logic programs under the answer set semantics, the description logics $\mathcal{SHOIN}(\mathbf{D})$ and $\mathcal{SHIF}(\mathbf{D})$, and Bayesian probabilities.

The work by Lutz and Schröder [14] introduces a family of probabilistic DLs the members of which relate to the probabilistic FOL of [8] in the same way as classical DLs relate to FOL. This family, denoted by Prob-DLs, introduces a set of probabilistic constructors $P_{\sim p}$ where $\sim \in \{<, \leq, =, \geq, >\}$ and $p \in [0, 1]$ to be applied to concepts and in some cases also to roles. If C is a concept, then $P_{\sim p}C$ is a concept which denotes objects that are an instance of C with probability $\sim p$. The semantics of Prob-DLs is based on probabilistic interpretations, which extend a classical DL interpretation with a probability distribution over a set of possible worlds. Concept subsumption in Prob-DLs refers to classical DL concept subsumption, although in every world. It is claimed that Prob-DLs are well-suited to capture aspects of uncertainty that are present in almost all biomedical ontologies. We believe, though, that they do not capture the intended semantics of ontology mappings.

7 Conclusions

Distributed ontology mappings are highly uncertain. We investigated the possibility to provide a reclassification semantics for weighted mappings extending DDL semantics. Reclassification semantics is based on the probability that individuals classified under a particular concept in one ontology would be classified in another concept in another ontology. Such a probabilistic view on weights should match the practice of matchers based on concept extensions or relation to a wider context, e.g., annotated resources.

We showed that such a semantics was preserving the classical DDL semantics in the sense that if crisp DDL mappings are encoded as weighted mappings with $[1, 1]$ weights, the consequences correspond. In fact, the reclassification semantics may be used as an alternative semantics to classical DDL semantics.

Then, weighted mapping entailment was defined from this semantics. Inferred mappings predict the probability of reclassification from weighted mappings. This allows

one to infer mappings across different mapping relations and term constructions. It also relates mapping inference to (crisp) ontological reasoning. As usual with probabilistic approaches, precision weakens with inference. We also discussed variation of the proposed framework.

There are several avenues for this work. The most direct one is to apply it to mapping (or ontology) debugging by ranking given and inferred mappings according to their weight intervals and help to detect those near mappings which would not appear as crisp mappings but are mappings of high weight.

Acknowledgements. Manuel Atencia is partially supported under Qualinca project (ANR-2012-CORD-012) sponsored by the French National Research Agency.

References

1. Atencia, M., Borgida, A., Euzenat, J., Ghidini, C., Serafini, L.: A formal semantics for weighted ontology mappings. Tech. Rep. 81401, Fondazione Bruno Kessler - IRST (2012)
2. Bao, J., Voutsadakis, G., Slutzki, G., Honavar, V.: Package-based description logics. In: Modular Ontologies: Concepts, Theories and Techniques for Knowledge Modularization, LNCS, vol. 5445, pp. 349–371. Springer (2009)
3. Borgida, A., Serafini, L.: Distributed description logics: Assimilating information from peer sources. J. Data Semantics 1, 153–184 (2003)
4. Calvanese, D.: Unrestricted and finite model reasoning in class-based representation formalisms. Number viii-96-2 of collana delle tesi del dottorato di ricerca in informatica, Dipartimento di Informatica e Sistemistica, Università di Roma "La Sapienza" (1996)
5. Euzenat, J., Shvaiko, P.: Ontology matching. Springer (2007)
6. Gal, A.: Uncertain Schema Matching. Synthesis Lectures on Data Management, Morgan & Claypool (2011)
7. Grau, B.C., Parsia, B., Sirin, E.: Combining owl ontologies using epsilon-connections. Journal of Web Semantics 4(1), 40–59 (2006)
8. Halpern, J.Y.: Reasoning about uncertainty. MIT Press (2003)
9. Heinsohn, J.: Probabilistic description logics. In: UAI '94: Proceedings of the 10th Annual Conference on Uncertainty in Artificial Intelligence. pp. 311–318. Morgan Kaufmann (1994)
10. Koller, D., Levy, A.Y., Pfeffer, A.: P-classic: A tractable probablistic description logic. In: Proceedings of the 14th National Conference on Artificial Intelligence (AAAI 97). pp. 390–397. AAAI Press (1997)
11. Lukasiewicz, T.: Expressive probabilistic description logics. Artificial Intelligence 172(6-7), 852–883 (2008)
12. Lukasiewicz, T., Predoiu, L., Stuckenschmidt, H.: Tightly integrated probabilistic description logic programs for representing ontology mappings. Annals of Mathematics and Artificial Intelligence 63(3-4), 385–425 (2011)
13. Lukasiewicz, T., Straccia, U.: Managing uncertainty and vagueness in description logics for the Semantic Web. Journal of Web Semantics 6(4), 291–308 (2008)
14. Lutz, C., Schröder, L.: Probabilistic description logics for subjective uncertainty. In: Proceedings of the 12th International Conference on the Principles of Knowledge Representation and Reasoning: (KR 2010). AAAI Press (2010)
15. Tournaire, R., Petit, J.M., Termier, A., Rousset, M.C.: Discovery of Probabilistic Mappings between Taxonomies: Principles and Experiments. J. Data Semantics 15, 66–101 (Apr 2011)
16. Zimmermann, A., Euzenat, J.: Three semantics for distributed systems and their relations with alignment composition. In: Proceedings of the 5th International Semantic Web Conference (ISWC 2006). LNCS, vol. 4273, pp. 16–29. Springer (2006)