Ideal Downward Refinement in the $\mathcal{EL}$ Description Logic

Keywords: Refinement Operators, Description Logics, OWL, EL, Inductive Logic Programming

Abstract

With the proliferation of the Semantic Web, there has been a rapidly rising interest in description logics, which form the logical foundation of the W3C standard ontology language OWL. While the number of OWL knowledge bases grows, there is an increasing demand for tools assisting knowledge engineers in building up and maintaining their structure. For this purpose, concept learning algorithms based on refinement operators have been investigated. In this paper, we provide an ideal refinement operator for the popular description logic $\mathcal{EL}$ and show that it is computationally feasible on large knowledge bases.

1. Introduction

The Semantic Web is steadily growing[1] and contains knowledge from diverse areas such as science, music, people, books, reviews, places, politics, products, software, social networks, as well as upper and general ontologies. The underlying technologies, sometimes called *Semantic Technologies*, are currently starting to create substantial industrial impact in application scenarios on and off the web, including knowledge management, expert systems, web services, e-commerce, e-collaboration, etc. Since 2004, the Web Ontology Language OWL, which is based on description logics (DLs), has been the W3C-recommended standard for Semantic Web knowledge representation and is a key to the growth of the Semantic Web.

However, recent progress in the field faces a lack of well-structured ontologies with large amounts of instance data due to the fact that engineering such ontologies constitutes a considerable investment of resources. Nowadays, knowledge bases often provide large amounts of instance data without sophisticated schemata. Methods for automated schema acquisition and maintenance are therefore being sought (see e.g. (Buitelaar et al., 2007)). In particular, concept learning methods have attracted interest, see e.g. (Esposito et al., 2004; Baader et al., 2007; Lehmann, 2007; Lehmann & Hitzler, 2008b).

Many concept learning methods borrow ideas from Inductive Logic Programming including the use of *refinement operators*. Properties like ideality, completeness, finiteness, properness, minimality, and non-redundancy are used as theoretical criteria for the suitability of such operators. It has been shown in (Lehmann & Hitzler, 2008a) that no ideal refinement operator for DLs such as $\mathcal{ALC}$, $\mathcal{SHOIN}$, and $\mathcal{SROIQ}$ can exist (the two latter DLs are underlying OWL and OWL 2, respectively). In this article, an important gap in the the analysis of refinement operator properties is closed by showing that ideal refinement operators for the DL $\mathcal{EL}$ do exist, which in turn can lead to a breakthrough in DL concept learning.

$\mathcal{EL}$ is a light-weight DL, but despite its limited expressive power it has proven to be of practical use in many real-world large-scale applications, e.g. the Systematized Nomenclature of Medicine Clinical Terms (SNOMED CT) (Bodenreider et al., 2007) and the GENE ONTOLOGY (The Gene Ontology Consortium, 2000). Since standard reasoning in $\mathcal{EL}$ is polynomial, it is suitable for large ontologies. It should furthermore be mentioned that $\mathcal{EL}^{++}$, an extension of $\mathcal{EL}$, will become one of three profiles in the upcoming standard ontology language OWL 2.

Overall, we make the following contributions:

- A gap in the research of properties of refinement operators in DLs is being closed.
- An ideal and practically useful refinement operator for $\mathcal{EL}$ is developed.
- Computational feasibility of the operator is shown.

In Section 2, we describe the preliminaries of our work and present the refinement operator in Section 3. We prove its ideality and describe how it can be optimised

[1] As a rough size estimate, the semantic index Sindice (http://sindice.com/) lists more than 10 billion entities from more than 100 million web pages.

to work efficiently and incorporate background knowledge. We evaluate the operator on real-world knowledge bases in Section 4. Related work is described in Section 5 and conclusions are drawn in Section 6.

2. Preliminaries

In this section, the definitions relevant for defining the refinement operator in Section 3 are being introduced. Besides recalling known facts from the literature, we introduce minimal $\mathcal{EL}$ trees that serve as the basis for the refinement operator.

2.1. The $\mathcal{EL}$ Description Logic

Before we begin to introduce the DL $\mathcal{EL}$, we briefly recall some notions from order theory. Let Q be a set and $\preceq$ a quasi order on Q, i.e., a reflexive and transitive binary relation on Q. Then $(Q, \preceq)$ is called a *quasi ordered space*. The quasi order $\preceq$ induces the equivalence relation $\simeq$ and the *strict quasi order* $\prec$ on Q: $q \simeq q'$ iff $q \preceq q'$ and $q' \preceq q$, and $q \prec q'$ iff $q \preceq q'$ and $q \not\simeq q'$. For $P \subseteq Q$, $sup(P) := \{p \in P \mid$ there is no $p' \in P$ width $p \prec p'\}$ defines the *supremum* of P. We say $(Q, \preceq)$ has a *greatest element* iff there is a $q^* \in Q$ such that $sup(Q) := \{q^*\}$.

The expressions in the DL $\mathcal{EL}$ are *concepts*, which are built inductively starting from sets of *concepts names* N_C and *role names* N_R of arbitrary but finite cardinality, and applying the *concept constructors* $\top$ (top), $C \sqcap D$ (conjunction) and $\exists r.C$ (existential restriction). By $\mathcal{C}(\mathcal{EL})$ we denote the set of all $\mathcal{EL}$ concepts. The semantics of an $\mathcal{EL}$ concept C is given in terms of an *interpretation* $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, where $\Delta^{\mathcal{I}}$ is a set called the *interpretation domain* and $\cdot^{\mathcal{I}}$ is the *interpretation function*. The interpretation function maps each $A \in N_C$ to a subset of $\Delta^{\mathcal{I}}$, and each $r \in N_R$ to a binary relation on $\Delta^{\mathcal{I}}$. It is then inductively extended to arbitrary $\mathcal{EL}$ concepts as $\top^{\mathcal{I}} := \Delta^{\mathcal{I}}$, $(C \sqcap D)^{\mathcal{I}} := C^{\mathcal{I}} \cap D^{\mathcal{I}}$ and $(\exists r.C)^{\mathcal{I}} := \{x \in \Delta^{\mathcal{I}} \mid$ there is $y \in C^{\mathcal{I}}$ with $(x, y) \in r^{\mathcal{I}}\}$. Subsequently, we will use A, B to identify concept names from N_C; C, D for $\mathcal{EL}$ concepts; and r, s for role names from N_R. By $|C|$ we denote the size of an $\mathcal{EL}$ concept, which is just the number of symbols used to write it down. When proving properties of $\mathcal{EL}$ concepts, the *role depth* of a concept C is a useful induction argument. It is defined by structural induction as $rdepth(A) = rdepth(\top) := 0$, $rdepth(C \sqcap D) := max(rdepth(C), rdepth(D))$ and $rdepth(\exists r.C) := rdepth(C)+1$. In this paper, an *($\mathcal{EL}$)-knowledge base* $\mathcal{K}$ is a finite union of *concept inclusion axioms* of the form $A \sqsubset B$, *role inclusion axioms* $r \sqsubset s$, *disjointness axioms* $A \sqcap B \equiv \bot$, *domain restriction axioms* $domain(r) = A$, and *range restriction axioms* $range(r) = A$. In the last two axioms, $A \in N_R$ and we assume w.l.o.g. that there is such an axiom for every $r \in N_R$. An interpretation $\mathcal{I}$ is a *model* of a knowledge base $\mathcal{K}$ iff for each $A \sqsubset B \in \mathcal{K}$, $A^{\mathcal{I}} \subset B^{\mathcal{I}}$; $r \sqsubset s \in \mathcal{K}$, $r^{\mathcal{I}} \subset s^{\mathcal{I}}$; $A \sqcap B \equiv \bot$, $A^{\mathcal{I}} \cap B^{\mathcal{I}} = \emptyset$; $domain(r) = A \in \mathcal{K}$, $x \in A^{\mathcal{I}}$ for all $(x, y) \in r^{\mathcal{I}}$; and for each $range(r) = A \in \mathcal{K}$, $y \in A^{\mathcal{I}}$ for all $(x, y) \in r^{\mathcal{I}}$. Given a knowledge base $\mathcal{K}$ and $\mathcal{EL}$ concepts C, D, C is *subsumed* by D w.r.t. $\mathcal{K}$ ($C \sqsubseteq_{\mathcal{K}} D$) iff $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$ for all models $\mathcal{I}$ of $\mathcal{K}$. In the remainder of this paper, we always assume a knowledge base to be implicitly present, and we therefore just write $C \sqsubseteq D$. It is important to mention that we can incorporate $\mathcal{L}$-knowledge bases $\mathcal{K}'$ that are defined using a DL $\mathcal{L}$ that is more expressive than $\mathcal{EL}$ by restricting $\mathcal{K}$ to be the $\mathcal{EL}$ fragment of $\mathcal{K}'$[2]. Moreover, we assume in the following that the $\mathcal{EL}$-knowledge base $\mathcal{K}$ is precomputed, and reasoning in $\mathcal{K}$ is in polynomial time. Obviously, $(\mathcal{C}(\mathcal{EL}), \sqsubseteq)$ forms a quasi ordered space, from which we can accordingly derive the relations $\equiv$ (*equivalence*) and $\sqsubset$ (*strict subsumption*). Given finite sets of concept names $\mathcal{A}, \mathcal{B} \subseteq N_C$, $\mathcal{A} \sqsubseteq \mathcal{B}$ iff for every $B \in \mathcal{B}$ there is some $A \in \mathcal{A}$ such that $A \sqsubseteq B$. We sometimes abuse notation and write $\mathcal{A} \sqsubseteq A$ instead of $\mathcal{A} \sqsubseteq \{A\}$. An $\mathcal{EL}$ concept C is satisfiable w.r.t. $\mathcal{K}$ iff there exists a model $\mathcal{I}$ of $\mathcal{K}$ such that $C^{\mathcal{I}} \neq \emptyset$.

Example 2.1 Let $N_C = \{\texttt{Human}, \texttt{Animal}, \texttt{Bird}, \texttt{Cat}\}$, $N_R = \{\texttt{has}, \texttt{has\_child}, \texttt{has\_pet}\}$, $\mathcal{K} = \{\texttt{has\_pet} \sqsubset \texttt{has}, \texttt{has\_child} \sqsubset \texttt{has}, \texttt{Bird} \sqsubset \texttt{Animal}, \texttt{Cat} \sqsubset \texttt{Animal}\}$, $C = \texttt{Human} \sqcap \exists \texttt{has}.\top$, and $D = \texttt{Human} \sqcap \exists \texttt{has\_child}.\texttt{Human} \sqcap \exists \texttt{has\_pet}.\texttt{Bird}$. Then we can infer $D \sqsubseteq C$. □

2.2. Downward Refinement Operators

Refinement operators are used to structure a search process for concepts. Intuitively, downward refinement operators construct specialisations of hypotheses. This idea is well-known in Inductive Logic Programming (Nienhuys-Cheng & de Wolf, 1997). Let $(Q, \preceq)$ be a quasi ordered space and denote by $\mathcal{P}(Q)$ the powerset of Q. A mapping $\rho : Q \to \mathcal{P}(Q)$ is a *downward refinement operator* on $(Q, \preceq)$ iff $q' \in \rho(q)$ implies $q' \preceq q$. In the remainder of this paper, we will call downward refinement operators just refinement operators. We write $q \rightsquigarrow_\rho q'$ for $q' \in \rho(q)$ and drop the index ρ if the refinement operator is clear from the context. A *refinement chain of length* n of a refinement operator ρ that starts in q_1 and ends in q_n is a sequence $q_1 \rightsquigarrow \ldots \rightsquigarrow q_n$ such that $q_i \rightsquigarrow q_{i+1}$ for $1 \leq i < n$. We say that the chain *goes through* q iff $q \in \{q_1, \ldots, q_n\}$. Moreover, $q \rightsquigarrow^* q'$ iff there exists a refinement chain

[2] By this restriction, we mean that every model of $\mathcal{K}'$ is also a model of $\mathcal{K}$, but not necessarily vice versa.

of length n starting from q and ending in q' for some $n \in \mathbb{N}$.

Refinement operators can be classified by means of their properties. Let $(Q, \preceq)$ be a quasi ordered space with a greatest element, and $q, q', q'' \in Q$. A refinement operator ρ is *finite* iff $\rho(q)$ is finite for any q. It is *proper* iff $q \rightsquigarrow q'$ implies $q \not\equiv q'$. We call ρ *complete* iff $q' \prec q$ implies $q \rightsquigarrow^* q''$ for some $q'' \equiv q'$. Let q^* be the greatest element in $(Q, \preceq)$, ρ is *weakly complete* iff for any $q' \prec q^*$, $q^* \rightsquigarrow^* q''$ with $q'' \equiv q'$. ρ is *redundant* iff $q^* \rightsquigarrow^* q'$ via two refinement chains, where one goes through a an element q'' and the other one does not go through q''. Finally, ρ is *ideal* iff it is finite, proper and complete.

2.3. Minimal $\mathcal{EL}$ Concepts

An important observation is that $\mathcal{EL}$ concepts can be viewed as directed labeled trees (Baader et al., 1999). This allows for deciding subsumption between concepts in terms of the existence of a simulation relation between the nodes of their corresponding trees, and moreover for a canonical representation of concepts as minimal $\mathcal{EL}$ trees.

An $\mathcal{EL}$ *graph* is a directed labeled graph $G = (V, E, \ell)$, where V is the finite set of *nodes*, $E \subseteq V \times N_R \times V$ is the set of *edges*, and $\ell : V \to \mathcal{P}(N_C)$ is the *labeling function*. We define $V(G) := V$, $E(G) := E$, $\ell(G) := \ell$ and $|G| := |V|$. For an edge $(v, r, w) \in E$, we call w an (r-)*successor* of v, and v an (r-)*predecessor* of w. Given a node $v \in V$, a labelling function ℓ and $L \subseteq N_C$, we define $\ell[v \to L]$ as $\ell[v \mapsto L](v) := L$ and $\ell[v \mapsto L](w) := \ell(w)$ for all $w \neq v$. Given G and $v \in V(G)$, we define $G[v \mapsto L] := (V(G), E(G), \ell(G)[v \mapsto L])$. We say $v_1 \xrightarrow{r_1} \ldots \xrightarrow{r_n} v_{n+1}$ is a *path* of length n from v_1 to v_{n+1} in G iff $(v_i, r_i, v_{i+1}) \in E$ for $1 \leq i \leq n$. A graph G contains a cycle iff there is a path $v \xrightarrow{r_1} \ldots \xrightarrow{r_n} v$ in G. A concept is represented by an $\mathcal{EL}$ *concept tree*, which is a connected finite $\mathcal{EL}$ graph t that does not contain any cycle, has a distinguished node called the *root* of t that has no predecessor, and every other node has exactly one predecessor along exactly one edge. The set of $\mathcal{EL}$ concept trees is denoted by T. In the following, we call an $\mathcal{EL}$ concept tree just a tree. Figure 3.1 illustrates some examples of such trees. Given a tree t, we denote by $root(t)$ its root. The tree t corresponding to a concept C is defined by induction on $n = rdepth(C)$. For $n = 0$, t consists of a single node that is labelled with all concepts names occurring in C. For $n > 0$, the root of t is labelled with all concept names occurring on the top-level of C. Furthermore, for each existential restriction $\exists r.D$ on the top-level of C, it has an r-labelled edge to the root of a subtree of t' which corresponds to D. As an example, the tree t corresponding to $A_1 \sqcap \exists r.A_2$ is $t = (\{v_1, v_2\}, \{(v_1, r, v_2)\}, \ell)$ where ℓ maps v_1 to $\{A_1\}$ and v_2 to $\{A_2\}$. By t_C we denote the tree corresponding to C. Obviously, the transformation from a concept to a tree can be performed in linear time w.r.t. the size of the concept. Similarly, any tree has a corresponding concept[3], and the transformation can be performed in linear time, too. We say t is *satisfiable* iff the concept C corresponding to t is satisfiable. Let t, t' be trees, $v \in V(t)$ and assume w.l.o.g. that $V(t) \cap V(t') = \emptyset$. Denote by $t[v \leftarrow (r, t)]$ the tree $(V(t) \cup V(t'), E(t) \cup E(t') \cup \{(v, r, root(t'))\}, \ell \cup \ell')$, where $\ell \cup \ell'$ is the obvious join of the labeling functions of t and t'. By $t(v)$ we denote the subtree at v. Let C be a concept and t the tree corresponding to C. We define $depth(t) := rdepth(C)$, and for $v \in V(t)$, $level(v) := depth(t) - depth(t(v))$. Moreover, $onlevel(t, n)$ is the set of nodes $\{v \mid level(v) = n\}$ that appear on level n.

Definition 2.2 Let $t = (V, E, \ell), t' = (V', E', \ell')$ be trees. A *simulation relation* from t' to t is a binary relation $\mathcal{S} \subseteq V \times V'$ such that $(root(t), root(t')) \in \mathcal{S}$ and if $(v, v') \in \mathcal{S}$ then the following *simulation conditions* are fulfilled:

(SC1) $\ell(v) \sqsubseteq \ell'(v')$

(SC2) for every $(v', r, w') \in E'$ there is $(v, r, w) \in E_1$ such that $r \sqsubseteq r'$ and $(w, w') \in \mathcal{S}$ □

We write $t \preceq t'$ if there exists a simulation relation $\mathcal{S}$ from t' to t. It is easily checked that $(T, \preceq)$ forms a quasi ordered space, and we derive the relations $\simeq$ and $\prec$ accordingly. A simulation $\mathcal{S}$ from t' to t is *maximal* if for every simulation $\mathcal{S}'$ from t' to t, $\mathcal{S}' \subseteq \mathcal{S}$. It is not hard to check that $\mathcal{S}$ is unique. Using a dynamic programming approach, it can be computed in $\mathcal{O}(|t| \cdot |t'|)$. The following lemma is proven by induction on $rdepth(D)$. It allows us to decide subsumption between concepts C, D in terms of the existence of a simulation between their corresponding trees t, t'.

Lemma 2.3 *Let C, D be concept with their corresponding trees t, t'. Then $C \sqsubseteq D$ iff $t \preceq t'$.*

The previous lemma allows us to interchange concepts and their corresponding trees, and the refinement operator presented in the next section will work on trees rather than concepts. We now introduce minimal $\mathcal{EL}$ trees which serve as a canonical representation of equivalent $\mathcal{EL}$ concepts. A similar topic, the minimization of XPath tree pattern queries, has been investigated in (Ramanan, 2002). In fact, $\mathcal{EL}$ trees generalize XPath tree pattern queries, and the relevant

[3] Strictly speaking, t has a set of corresponding concepts, which are all equivalent up to commutativity.

algorithms carry over straightforwardly to $\mathcal{EL}$ trees.

Definition 2.4 Let $t = (V, E, \ell)$ be a tree. We call t *label reduced* if for all $v \in V$ there does not exist $\mathcal{B} \subseteq N_C$ such that $|\mathcal{B}| < |\ell(v)|$ and $\mathcal{A} \equiv \mathcal{B}$. Moreover, t *contains redundant subtrees* if there are $(v, r, w), (v, r', w') \in E$ with $w \neq w'$, $r \sqsubseteq r'$ and $t(w) \preceq t(w')$. We call t *minimal* if t is label reduced and does not contain redundant subtrees. □

It follows that minimality of a tree t can be checked in $\mathcal{O}(|t|^2)$ by computing the maximal simulation from t to t and then checking for each $v \in V(t)$ whether v is label reduced and, using $\mathcal{S}$, whether v is not the root of redundant subtrees. The set of minimal $\mathcal{EL}$ trees is denoted by T_{min}. We close this section with two lemmas that will be helpful in the next section.

Lemma 2.5 *Let T_n be the set of minimal $\mathcal{EL}$ trees up to depth $n \geq 0$. Then $|T_n|$ is finite.*

Lemma 2.6 *Given trees t, t' with $depth(t') < depth(t)$. Then $t' \not\preceq t$.*

3. An Ideal $\mathcal{EL}$ Refinement Operator

In this section, we define an ideal refinement operator. In the first part, we are more concerned with a description of the operator on an abstract level, which allows us to prove its properties. The next part addresses optimisations of the operator that make it more usable in practice.

3.1. Definition of the Operator

For simplicity, we subsequently assume the knowledge base to only contain concept and role inclusion axioms. We will sketch in the next section how the remaining restriction axioms can be incorporated in the refinement operator.

The refinement operator ρ, to be defined below, is a function that maps a tree $t \in T_{min}$ to a subset of T_{min}. It can be divided into the three base operations *label extension*, *label refinement* and *edge refinement*. Building up on that, the complex operation *attach subtree* is defined. Each such operation takes a tree $t \in T_{min}$ and a node $v \in V(t)$ as input and returns a set of trees that are refined at node v. In order to provide a definition of the base operations, we introduce the function $sh_\downarrow$ which —roughly speaking— allows us to "climb down" the subsumption hierarchy under some constraints. Let $\mathcal{A} \subseteq N_C$ and $\mathcal{B} \subseteq N_C \cup \{\top\}$, $sh_\downarrow(\mathcal{A}, \mathcal{B}) = sup\{A \in N_C \mid \mathcal{A} \not\sqsubseteq A$ and there is exactly one $B \in \mathcal{B}$ such that $A \sqsubset B\}$. For example, let $N_C = \{A_1, A_2, A_3, A_4\}$ and $\mathcal{K} = \{A_2 \sqsubset A_1, A_3 \sqsubset A_1\}$. Then $sh_\downarrow(\{A_2\}, \{\top\}) = \{A_3, A_4\}$ and $sh_\downarrow(\emptyset, \{A_1, A_4\}) = \{A_2, A_3\}$. Likewise, we define $sh_\downarrow$ on role names from N_R: $sh_\downarrow(r) := sup\{s \in N_R \mid s \sqsubset r\}$. The base operations are as follows: the operation $e\ell(t, v)$ returns the set of those minimal satisfiable trees that are derived from t by extending the label of v. Likewise, $r\ell(t, v)$ is the set of trees obtained from t by refining the label of v. Last, $re(t, v)$ is obtained from t by refining any of the outgoing edges at v. Formally,

- $e\ell(t, v)$: $t' \in e\ell(t, v)$ iff $t' \in T_{min}$ and

  $t' = t[v \mapsto \ell(v) \cup \{A\}]$ for $A \in sh_\downarrow(\ell(v), \{\top\})\}$

- $r\ell(t, v)$: $t \in r\ell(t, v)$ iff $t' \in T_{min}$ and

  $t' = t[v \mapsto (\ell(v) \cup \{A\}) \setminus sh_\uparrow(A)]$ for $A \in sh_\downarrow(\emptyset, \ell(v))$

- $re(t, v)$:$t' \in re(t, v)$ iff $t' \in T_{min}$ and $t' = (V, E', \ell)$, where $E' = E \setminus \{(v, r, w)\} \cup \{(v, r', w)\}$ for some $(v, r, w) \in E$ and $r' \in sh_\downarrow(r)$

Algorithm 1 Computation of the set $as(t, v)$

$\mathcal{T} := \emptyset$; $\mathcal{M} := \{(t_\top, N_R)\}$;
while $\mathcal{M} \neq \emptyset$ **do**
  choose and remove $(t', \mathcal{R}) \in \mathcal{M}$;
  $\mathcal{R}' := sup(\mathcal{R})$; $\mathcal{R}'' := \emptyset$;
  while $\mathcal{R}' \neq \emptyset$ **do**
    choose and remove $r \in \mathcal{R}'$;
    $t'' := t[v \leftarrow (r, t')]$; $w := root(t')$;
    if t'' is minimal **then**
      $\mathcal{T} := \mathcal{T} \cup \{t''\}$;
    else
      for all $(v, r', w') \in E(t'')$ with $w \neq w'$ and $r \sqsubseteq r'$ **do**
        if $t''(w) \preceq t''(w')$ **then**
          nextwhile;
        end if
      end for
      $\mathcal{R}' := \mathcal{R}' \cup (sh_\downarrow(r) \cap \mathcal{R})$; $\mathcal{R}'' := \mathcal{R}'' \cup \{r\}$;
    end if
  end while
  $\mathcal{M} := \mathcal{M} \cup \{(t^*, \mathcal{R}'') \mid t^* \in \rho(t'), \mathcal{R}'' \neq \emptyset\}$;
end while
return $\mathcal{T}$;

The crucial part of the refinement operator is the attach subtree operation, which is defined by Algorithm 1. The set $as(t, v)$ consists of minimal trees obtained from t that have an extra subtree attached to v. It recursively calls the refinement operator ρ and we therefore give its definition before we explain $as(t, v)$ in more detail.

Definition 3.1 The refinement operator $\rho : T_{min} \rightarrow \mathcal{P}(T_{min})$ is defined as:

$$\rho(t) := \bigcup_{v \in V(t)} (e\ell(t,v) \cup r\ell(t,v) \cup re(t,v) \cup as(t,v))$$

New edges for an input tree are introduced by as. For $t \in T_{min}$ and $v \in V$, $as(t,v)$ keeps a set of output trees $\mathcal{T}$ and a set $\mathcal{M}$ of candidates which are tuples consisting of a minimal satisfiable $\mathcal{EL}$ tree and a set of role names. Within the first while loop, an element $(t', \mathcal{R})$ is removed from $\mathcal{M}$. The set $\mathcal{R}'$ is initialized to contain the greatest elements of $\mathcal{R}$, and $\mathcal{R}''$ is initially empty and will later contain role names that need further inspection. In the second while loop, the algorithm iterates over all role names r in $\mathcal{R}'$. First, the tree t'' is constructed from t by attaching the subtree (v,r,w) to v, where w is the root of t'. It is then checked whether t'' is minimal. If this is the case, t'' is a refinement of t and is added to $\mathcal{T}$. Otherwise there are two reasons why t'' is not minimal: Either the newly attached subtree is subsumed by some other subtree of t, or the newly attached subtree subsumes some other subtree of t. The latter case is checked in Line 11, and if it applies the algorithm skips the loop. This prevents the algorithm from running into an infinite loop, since we would not be able to refine t' until t'' becomes a minimal tree. Otherwise in the former case, we proceed in two directions. First, $sh_{\downarrow}(r)$ is added to $\mathcal{R}'$, so it can be checked in the next round of the second while loop whether t' attached via some $r' \in sh_{\downarrow}(r) \cap \mathcal{R}$ to v yields a refinement. Second, we add r to $\mathcal{R}''$, which can be seen as "remembering" that r did not yield a refinement in connection with t'. Finally, once $\mathcal{R}'$ is empty, in Line 19 we add all tupels $(t^*, \mathcal{R}'')$ to $\mathcal{M}$, where t^* is obtained by recursively calling ρ on t'.

Example 3.2 Figure 3.1 depicts the set $\rho(\texttt{Human} \sqcap \exists\texttt{has.Animal})$ w.r.t. $\mathcal{K}$ from Example 2.1. □

Proposition 3.3 *ρ is a is a finite, proper and weakly complete downward refinement operator on $(T_{min}, \preceq)$.*

Proof In the following, let $t \in T_{min}$ and $v \in V(t)$.

First, it is easily seen that ρ is a downward refinement operator. Every operation of ρ adds a label or a subtree to a node v, or replaces a label or edge-label by a refined label or edge respectively. Hence, $t' \preceq t$ for all $t' \in \rho(t)$.

Regarding finiteness of ρ, Lemma 2.5 guarantees that there is only a finite number of minimal $\mathcal{EL}$ trees up to a fixed depth. It follows from Lemma 2.6 that for a given tree t, $\rho(t)$ only consists of trees of depth at most $depth(t) + 1$. Hence, $\rho(t)$ is finite.

In order to prove properness of ρ, it is sufficient to show $t \not\preceq t'$ for $t' \in \rho(t)$. To the contrary, assume $t \preceq t'$ and that t has been refined at v. Let $\mathcal{S}$ be a simulation from t' to t. Since v has been refined, it is not hard to show that $(v,v) \notin \mathcal{S}$. We have that $\mathcal{S}$ is a simulation, so there must be some $v' \in V(t)$ with $level(v') = level(v)$ such that $(v', v) \in \mathcal{S}$. This implies that there is a simulation $\mathcal{S}'$ on t' with $\{(v',v),(v,v)\} \subseteq \mathcal{S}'$. It follows that t' contains a redundant subtree at the predecessor of v, contradicting to the minimality of t'.

Regarding weakly completeness, let $depth(t) \leq n$. We show that t is reachable from $t_\top$ by nested induction on n and $m := |\{(root(t), r, w) \in E(t)\}|$. For the induction base case $n = 0, m = 0$, t is just a single node labeled with some concept names. It is easily seen that by repeatedly applying $e\ell(t,v)$ and $r\ell(t,v)$ to this node we eventually reach t. For the induction step, let $n > 0, m > 0$. Hence, t is a tree with m successor nodes $w_1, \ldots, w_m$ attached along edges $r_1, \ldots, r_m$ to t. By the induction hypothesis, the tree t_{m-1}, which is obtained from t by removing the subtree $t(w_1)$ from t, is reachable from $t_\top$ and denote by θ the corresponding refinement chain. Consequently, we have $t' = t_{m-1}[v \leftarrow (r', t'_{w_1})] \in as(t_m, v)$ for some tree t'_{w_1} occurring in θ and $r_1 \sqsubseteq r'_1$. Note that no intermediate tree in the refinement chain from $t_\top$ to t'_1 is dropped in Line 13 of as due to the minimality of t and the fact that $t(w_1) \preceq t'_{w_1}$. Now by first repeatedly applying the remaining refinement steps from θ to t' and then repeatedly refining r'_1 with $re(t,v)$, we reach a tree t_m. ■

Still, ρ is not ideal, since it is not complete. It is however easy to derive a complete operator ρ^* from ρ:

$$\rho^*(t) \;\; := \;\; sup\{t' \mid t_\top \rightsquigarrow^*_\rho t', t' \prec t \text{ and } depth(t') \leq depth(t) + 1\}.$$

This construction is needed, because we would for example not be able to reach $\exists r.(A_1 \sqcap A_2)$ starting from $\exists r.A_1 \sqcap \exists r.A_2$ with ρ.

Theorem 3.4 *The $\mathcal{EL}$ downward refinement operator ρ^* is ideal.*

In (Lehmann & Hitzler, 2008a) it has been shown for languages other than $\mathcal{EL}$ that complete and non-redundant refinement operators do not exist under a mild assumption. It is also not hard to show that ρ has to be redundant in our setting:

Proposition 3.5 *Let $\psi : T_{min} \rightarrow \mathcal{P}(T_{min})$ be a complete refinement operator. Then ψ is redundant.*

3.2. Optimisations

We used two different kinds of optimisations: The first is concerned with the performance of minimality tests

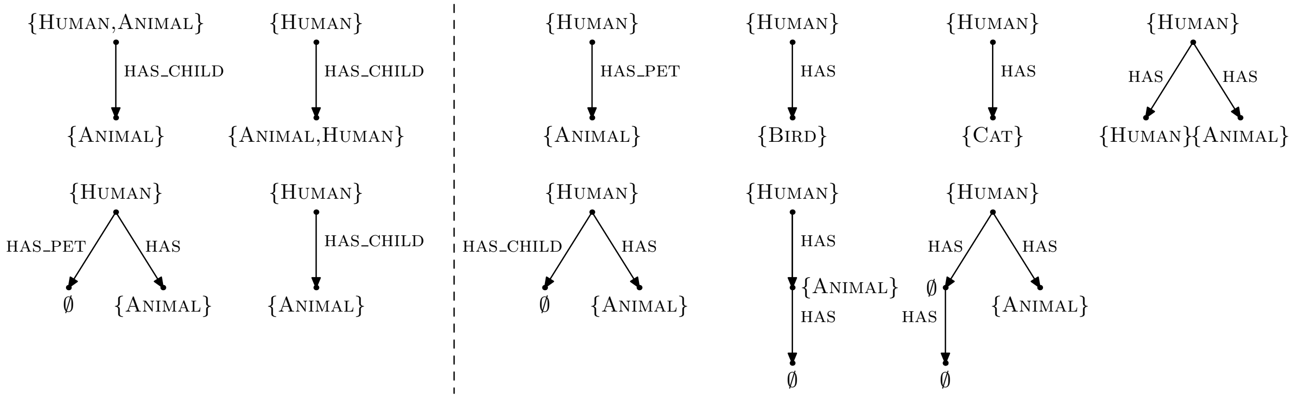

Figure 1. The set $\rho(\texttt{Human} \sqcap \exists\texttt{has.Animal})$ of minimal satisfiable trees w.r.t. the knowledge-base $\mathcal{K}$ from Example 2.1.

and the second reduces the number of trees returned by ρ by incorporating more background knowledge.

Recall from Section 2.3 that checking for minimality of a tree t involves computing a maximal simulation $\mathcal{S}$ on $V(t)$ and is in $\mathcal{O}(|t|^2)$. In order to avoid expensive re-computations of $\mathcal{S}$ after each refinement step, the data-structure of t is extended such that sets $\mathcal{C}_1^{\leftarrow}(v)$, $\mathcal{C}_1^{\rightarrow}(v)$, $\mathcal{C}_2^{\leftarrow}(v)$ and $\mathcal{C}_2^{\rightarrow}(v)$ are attached to every node $v \in V(t)$. Here, the set $\mathcal{C}_1^{\leftarrow}(v)$ contains those nodes w such that (SC1) holds for (v, w) according to Definition 2.2. Likewise, $\mathcal{C}_2^{\rightarrow}(v)$ is the set of those nodes w such that (SC2) holds for (w, v), and $\mathcal{C}_1^{\leftarrow}(v)$ and $\mathcal{C}_2^{\rightarrow}(w)$ are defined accordingly. When checking for minimality, it is moreover sufficient that each such set is restricted to only consist of nodes from *onlevel*(v) excluding v itself. This fragmentation of $\mathcal{S}$ allows us to perform local updates instead of re-computation of $\mathcal{S}$ after an operation is performed on v. For example, when the label of v is extended, we only need to recompute $\mathcal{C}_1^{\leftarrow}(v)$, update $\mathcal{C}_1^{\rightarrow}(w)$ for every $w \in \mathcal{C}_1^{\leftarrow}(v)$, and then repeatedly update $\mathcal{C}_2^{\rightarrow}(v')$ and $\mathcal{C}_2^{\leftarrow}(v')$ for every predecessor node v' of an updated node until we reach the root of t. This method saves a considerable amount of computation, since the number of nodes affected by an operation is empirically relatively small.

In order to keep $|\rho(t)|$ small, we use role domains and ranges as well as disjoint concepts inferred from $\mathcal{K}$. The domain restriction axioms can be used to reduce the set of role names considered when adding a subtree or refining an edge: For instance, let w be a node, (v, r, w) the edge pointing to w, and *range*$(r) = A$. When adding an edge (w, s, u), we ensure that *range*$(r) \sqcap$ *domain*(s) is satisfiable. This ensures that only compatible roles are combined. Similar effects are achieved by mode declarations in ILP tools. However, in OWL ontologies role domains and ranges are usually already present and hence do not need to be added manually. Similar optimisations can be applied to edge refinement. In *as*(t, v), we furthermore use range restrictions to automatically label a new node with the corresponding role range. For example, if the edge has label r and *range*$(r) = A$, then the new node w is assigned label $\ell(w) = \{A\}$ (instead of $\ell(w) = \emptyset$).

We now address the optimisation of extending node labels (function $e\ell$). Let A be a concept name for which we want to know whether to add it to $\ell(v)$. We first check $A \sqsubseteq \ell(v)$. If yes, we discard A since we could reach an equivalent concept by refining a concept in $\ell(v)$, i.e. we perform redundancy reduction. Let (u, r, v) be the edge pointing to v and *range*$(r) = B$. We verify that $A \sqcap B$ is satisfiable and discard A otherwise. Additionally as before, we test whether $\ell(v) \sqsubseteq A$. If yes, then A is also discarded, because adding it would not result in a proper refinement. Performing the last step in a top down manner, i.e. start with the most general concepts A in the class hierarchy, ensures that we compute the supremum of eligible concepts, which can be added to $\ell(v)$. In summary, we make sure that the tree we obtain is label reduced, and perform an on-the-fly test for satisfiability. Applying similar ideas to the case of label refinement is straight forward. In practice, the techniques described here narrow the set of trees returned in a refinement step significantly by ruling out concepts, which are unsatisfiable w.r.t. $\mathcal{K}$ or which can also be reached via other refinement chains. This is is illustrated by the following example.

Example 3.6 Let $\mathcal{K}$ be as in Example 2.1 and define $\mathcal{K}' := \mathcal{K} \cup \{$*domain*$(\texttt{has\_pet}) = \texttt{Animal}$, *domain*$(\texttt{has\_child}) = \texttt{Human}$, *range*$(\texttt{has\_child}) = \texttt{Human}$, $\texttt{Human} \sqcap \texttt{Animal} \equiv \bot\}$. By incorporating the additional axioms, $\rho(\texttt{Human} \sqcap \exists\texttt{has.Animal})$ only contains the trees on the right-hand side of the dashed line in Figure 3.1, except for $\texttt{Human} \sqcap \exists\texttt{has\_child}.\top \sqcap \exists\texttt{has.Animal}$, which becomes $\texttt{Human} \sqcap \exists\texttt{has\_child.Human} \sqcap \exists\texttt{has.Animal}$ due to the range of $\texttt{has\_child}$. □

name	logical axioms	classes	roles	ρ av. time (in ms)	ρ per ref. (in ms)	reasoning time (%)	refinements av.	refinements max.	ref. size av.	ref. size max.
Genes	42656	26225	4	167.2	0.14	68.4	1161.5	2317	5.0	8
CTON	33203	17033	43	76.2	0.08	5.1	220.2	28761	5.8	24
Galen	4940	2748	413	3.5	0.21	37.1	17.0	346	4.9	16
Process	2578	1537	102	193.6	0.16	27.2	986.5	23012	5.7	22
Transport	1157	445	89	164.4	0.09	5.9	985.2	22651	5.7	24
Earthrealm	931	559	81	407.4	0.17	23.2	1710.3	27163	5.7	19
Tambis	595	395	100	141.6	0.09	1.5	642.4	26685	5.8	23

Table 1. Benchmark results on ontologies from the TONES repository. The results show that ρ works well even on large knowledge bases. The time needed to compute a refinement is below one millisecond and does not show large variations.

4. Evaluation of the Operator

To evaluate the operator, we computed *random refinement chains* of ρ. A random refinement chain is obtained by applying ρ to $\top$, choosing one of the refinements uniformly at random, then applying ρ to this refinement etc.

To asses the performance of the operator, we tested it on real ontologies chosen from the TONES repository[4], including some of the most complex OWL ontologies. We generated 100 random refinement chains of length 8 and measured the results. We found experimentally that this allows us to evaluate the refinement operator on a diverse set of concept trees. The tests were run on an Athlon XP 4200+ (dual core 2.2 GHz) with 4 GB RAM. As a reasoner we used Pellet 1.5. The benchmarks do not include the time to load the ontology into the reasoner and classify it.

The results are shown in Table 1. The first four columns contain the name and relevant statistics of the ontology considered. The next column shows the average time the operator needed on each input concept. In the following column this value is divided by the number of refinements of the input concept. The subsequent column shows how much time is spend on reasoning during the computation of refinements. The two last columns contain the number of refinements obtained and their size. Here, we measure size as the number of nodes in a concept tree plus the sum of the cardinality of all node labels.

The most interesting insight from Table 1 is that despite the different size and complexity of the ontologies, the time needed to compute a refinement is low and does not show large variations (between 0.09 and 0.21 ms). This indicates that the operator scales well to large knowledge bases. It can also be observed that the number of refinements can be very high in certain cases, which is due to the large number of classes and properties in many ontologies and the absence of explicit or implicit disjointness between classes. We want to note that when the operator is used to learn concepts from instances (standard learning task), one can use the optimisations in Section 3.2 and consider classes without common instances instead of class disjointness. In this case, the number of refinements of a given concept will usually be much lower, since no explicit disjointness axioms are required. In all experiments we also note that the time the reasoner requires differs a lot (from 1.5% to 68.4%). However, since the number of reasoner requests is finite and the results are cached, this ratio will decrease with more calls to the refinement operator. Summing up, the results show that efficient ideal refinement on large ontologies can be achieved in $\mathcal{EL}$, which in turn is promising for $\mathcal{EL}$ concept learning algorithms.

[4]http://owl.cs.manchester.ac.uk/repository/

5. Related Work

In the area of Inductive Logic Programming considerable efforts have been made to analyse the properties of refinement operators (for a comprehensive treatment, see e.g. (Nienhuys-Cheng & de Wolf, 1997)). The investigated operators are usually based on horn clauses. In general, applying such operators to DL problems is considered not be a good choice (Badea & Nienhuys-Cheng, 2000). However, some of the theoretical foundations of refinement operators in Horn logics also apply to description logics, which is why we want to mention work in this area here.

In Shapiro's Model Inference System (Shapiro, 1991), he describes how refinement operators can be used to adapt a hypothesis to a sequence of examples. In the following years, refinement operators became widely used. (van der Laag & Nienhuys-Cheng, 1994) found some general properties of refinement operators in quasi-ordered spaces. Nonexistence conditions for ideal refinement operators relating to infinite ascending and descending refinement chains and covers have been developed. The result has been used to show the non-existence of ideal refinement operators for clauses ordered by θ-subsumption. Later, refinement operators have been extended to theories (clause sets) (Fanizzi et al., 2003).

Within the last decade, several refinement operators for description logics have been investigated. The most fundamental work is (Lehmann & Hitzler, 2008a), which shows for many description languages the maximal sets of properties which can be combined. Among other things, a non-ideality result for the languages $\mathcal{ALC}$, $\mathcal{SHOIN}$, and $\mathcal{SROIQ}$ is shown. We extend this work by providing an ideality result for $\mathcal{EL}$. Refinement operators for $\mathcal{ALER}$ (Badea & Nienhuys-Cheng, 2000), $\mathcal{ALN}$ (Fanizzi et al., 2004), $\mathcal{ALC}$ (Lehmann & Hitzler, 2008b; Iannone et al., 2007) have been created and used in learning algorithms. (Esposito et al., 2004) and (Fanizzi et al., 2004) have stated that further research into refinement operator properties is required for building the theoretical foundations of learning in DLs. Finally, (Lisi & Malerba, 2003) provides ideal refinement in $\mathcal{AL}$-log, a hybrid language merging Datalog and $\mathcal{ALC}$, but naturally a different order than DL subsumption was used.

6. Conclusions and Future Work

In summary, we have provided an efficient ideal $\mathcal{EL}$ refinement operator, thereby closing a gap in refinement operator research. We have shown that the operator can be applied to very large ontologies and makes profound use of background knowledge. In future work, we want to incorporate the refinement operator in learning algorithms, and investigate whether certain extensions of $\mathcal{EL}$ may be supported by the operator without losing ideality.

References

Baader, F., Molitor, R., & Tobies, S. (1999). Tractable and decidable fragments of conceptual graphs. *Seventh International Conference on Conceptual Structures (ICCS'99)* (pp. 480–493). Springer Verlag.

Baader, F., Sertkaya, B., & Turhan, A.-Y. (2007). Computing the least common subsumer w.r.t. a background terminology. *J. Applied Logic*, *5*, 392–420.

Badea, L., & Nienhuys-Cheng, S.-H. (2000). A refinement operator for description logics. *Proceedings of the 10th International Conference on Inductive Logic Programming* (pp. 40–59). Springer-Verlag.

Bodenreider, O., Smith, B., Kumar, A., & Burgun, A. (2007). Investigating subsumption in SNOMED CT: An exploration into large description logic-based biomedical terminologies. *Artificial Intelligence in Medicine*, *39*, 183–195.

Buitelaar, P., Cimiano, P., & Magnini, B. (Eds.). (2007). *Ontology learning from text: Methods, evaluation and applications*, vol. 123 of *Frontiers in Artificial Intelligence*. IOS Press.

Esposito, F., Fanizzi, N., Iannone, L., Palmisano, I., & Semeraro, G. (2004). Knowledge-intensive induction of terminologies from metadata. *Third International Semantic Web Conference* (pp. 441–455). Springer.

Fanizzi, N., Ferilli, S., Iannone, L., Palmisano, I., & Semeraro, G. (2004). Downward refinement in the ALN description logic. *HIS* (pp. 68–73). IEEE Computer Society.

Fanizzi, N., Ferilli, S., Mauro, N. D., & Basile, T. M. A. (2003). Spaces of theories with ideal refinement operators. *Proc. of 18th Int. Joint Conf. on Artificial Intelligence* (pp. 527–532). Morgan Kaufmann.

Iannone, L., Palmisano, I., & Fanizzi, N. (2007). An algorithm based on counterfactuals for concept learning in the semantic web. *Applied Intelligence*, *26*, 139–159.

Lehmann, J. (2007). Hybrid learning of ontology classes. *Machine Learning and Data Mining in Pattern Recognition, 5th International Conference* (pp. 883–898). Springer.

Lehmann, J., & Hitzler, P. (2008a). Foundations of refinement operators for description logics. *Proc. of 17th Int. Conf. on Inductive Logic Programming (ILP 2007)* (pp. 161–174). Springer. Best Student Paper.

Lehmann, J., & Hitzler, P. (2008b). A refinement operator based learning algorithm for the ALC description logic. *Proc. of 17th Int. Conf. on Inductive Logic Programming (ILP 2007)* (pp. 147–160). Springer. Best Student Paper.

Lisi, F. A., & Malerba, D. (2003). Ideal refinement of descriptions in AL-log. *Proc. of 13th Int. Conf. on Inductive Logic Programming (ILP)* (pp. 215–232). Springer.

Nienhuys-Cheng, S.-H., & de Wolf, R. (Eds.). (1997). *Foundations of inductive logic programming.* Lecture Notes in Computer Science. Springer.

Ramanan, P. (2002). Efficient algorithms for minimizing tree pattern queries. *SIGMOD '02: Proc. of the 2002 ACM SIGMOD Int. Conf. on Management of data* (pp. 299–309). ACM.

Shapiro, E. Y. (1991). Inductive inference of theories from facts. In J. L. Lassez and G. D. Plotkin (Eds.), *Computational logic: Essays in honor of Alan Robinson*, 199–255. The MIT Press.

The Gene Ontology Consortium (2000). Gene ontology: tool for the unification of biology. *Nature Genetics*, *25*, 25–29.

van der Laag, P. R. J., & Nienhuys-Cheng, S.-H. (1994). Existence and nonexistence of complete refinement operators. *Proc. of 7th Europ. Conf. on Machine Learning* (pp. 307–322). Springer-Verlag.

A. Appendix

A.1. Omitted Proofs

Definition A.1 Let C be a concept. The tree t corresponding to C is defined by induction on $n = rdepth(C)$. For $n = 0$, we have that $C = A_1 \sqcap \ldots \sqcap A_k$ and define $t := (\{v\}, \emptyset, \ell)$ with $\ell(v) := \{A_1, \ldots, A_k\}$. For $n > 0$, $C = A_1 \sqcap \ldots \sqcap A_k \sqcap \exists r_1.D_1 \sqcap \ldots \sqcap \exists r_m.D_m$ with $rdepth(D_i) < n, 1 \leq i \leq m$. By the induction hypothesis, for each D_i there exists a tree $t_i = (V_i, E_i, \ell_i), 1 \leq i \leq m$. Without loss of generality assume $V_i \cap V_j = \emptyset, 1 \leq i \neq j \leq m$. Define $t := (V, E, \ell)$ where

- $V := \{v\} \cup \bigcup_{1 \leq i \leq m} V_i$
- $E := \{(v, r_i, root(t_i)) \mid 1 \leq i \leq m\} \cup \bigcup_{1 \leq i \leq m} E_i$
- $$\ell(w) := \begin{cases} \{A_1, \ldots, A_k\} & \text{if } w = v \\ \ell_i(w) & \text{if } w \in V_i \end{cases}$$

Definition A.2 Let $\mathcal{I}$ be an interpretation. The $\mathcal{EL}$ graph $\mathcal{G}_\mathcal{I} = (V_\mathcal{I}, E_\mathcal{I}, \ell_\mathcal{I})$ corresponding to $\mathcal{I}$ is defined as follows:

- $x \in V_\mathcal{I}$ iff $x \in \Delta^\mathcal{I}$
- $(x, r, y) \in E_\mathcal{I}$ iff $(x, y) \in r^\mathcal{I}$
- $A \in \ell_\mathcal{I}(x)$ iff $x \in A^\mathcal{I}$ □

This definition also allows us to view $\mathcal{EL}$ graphs and in particular $\mathcal{EL}$ trees as interpretations.

Lemma A.3 *Let C be an $\mathcal{EL}$ concept with the corresponding $\mathcal{EL}$ tree $t = (V, E, \ell)$ with root v, and let $\mathcal{I}$ be an interpretation with $x \in \Delta^\mathcal{I}$ and the corresponding $\mathcal{EL}$ graph $\mathcal{G} = (V_\mathcal{I}, E_\mathcal{I}, \ell_\mathcal{I})$. Then $x \in C^\mathcal{I}$ iff there exists a simulation $\mathcal{S}$ from v to x.*

Proof The proof is by induction on $d = rdepth(C)$ in both directions.

($\Rightarrow$) For the induction base case, let $d = 0$ and $x \in (A_1 \sqcap \ldots \sqcap A_k)^\mathcal{I}$. Define $\mathcal{S} := \{(x, v)\}$, which obviously is a simulation. Now for the induction step, let $x \in (A_1 \sqcap \ldots \sqcap A_k \sqcap \exists r_1.C_1 \sqcap \ldots \sqcap \exists r_m.C_m)^\mathcal{I}$. There are $(x, v_i) \in r_i^\mathcal{I}$ such that $v_i \in C_i^\mathcal{I}$, $(x_C, r_i, v_i) \in E_C$ and by the induction hypothesis there exist simulations $\mathcal{S}_i$ from v_i to x_i for $1 \leq i \leq m$. Hence, $\mathcal{S} := \bigcup_{1 \leq i \leq m} \mathcal{S}_i \cup \{(x, v)\}$ is a simulation from v to x.

($\Leftarrow$) For the induction base case, let $d = 0$ and $C = A_1 \sqcap \ldots \sqcap A_k$. For every $A \in \ell(v)$ we have $A' \in \ell_\mathcal{I}(x)$ with $A' \sqsubseteq A$, so clearly $x \in C^\mathcal{I}$. For the induction step, let $C = A_1 \sqcap \ldots \sqcap A_k \sqcap \exists r_1.C_1 \sqcap \ldots \sqcap \exists r_m.C_m$ and $\mathcal{S}$ be a simulation from v to x. There are $(v, r_i, v_i) \in E$ and by the simulation conditions, there are also $(x, x_i) \in r_i^\mathcal{I}$ and $(x_i, v_i) \in \mathcal{S}$ for $1 \leq i \leq m$. Now $\mathcal{S}$ is a simulation from each v_i to x_i, and hence by the induction hypothesis $x_i \in C_i^\mathcal{I}$. Consequently, $x \in C^\mathcal{I}$. ■

Lemma 2.3 Let C, D be concept with their corresponding trees t, t'. Then $C \sqsubseteq D$ iff $t \preceq t'$.

Proof In the following let $x_C = root(t)$ and $x_D = root(t')$.

($\Rightarrow$) We show the contrapositive. Assume there does not exist a simulation from x_D to x_C. Now the identity on the vertices of t_C is a simulation from x_C to x_C and hence $x_C \in C^\mathcal{I}$, where $\mathcal{I}$ is the interpretation corresponding to t_C. Since there does not exist a simulation from x_D to x_C, the previous lemma gives $x_C \notin D^\mathcal{I}$.

($\Leftarrow$) Let $\mathcal{S}$ be a simulation from x_D to x_C, and let $\mathcal{I}$ be an interpretation with $y \in C^\mathcal{I}$. By the previous lemma, there exists a simulation $\mathcal{S}'$ and the composition $\mathcal{S} \circ \mathcal{S}'$ yields a simulation from x_D to y. Hence $y \in D^\mathcal{I}$. ■

Lemma A.4 *Given trees t_1, t_2, the maximal simulation from t_2 to t_1 is unique.*

Proof Suppose there are maximal simulations S_1, S_2 from t_2 to t_1 with $S_1 \not\subseteq S_2$ and $S_2 \not\subseteq S_1$. Observe that $S := S_1 \cup S_2$ is a simulation from t_2 to t_1, $S_1 \subset S$ and $S_2 \subset S$, which contradicts to the maximality of S_1 and S_2. ■

Lemma 2.5 Let T_n be the set of minimal $\mathcal{EL}$ trees up to depth $n \geq 0$. Then $|T_n|$ is finite.

Proof The proof is by induction on $n = 0$. We have $T_0 = 2^{|N_C|}$. For the induction step, assume T_{n+1} is infinite. Hence, there is a tree $t \in T_{n+1}$ whose root v has more than $|N_R| \cdot |T_n|$ outgoing edges. Consequently, there are distinct $(v, r, w), (v, r, w') \in E$ such that $t(w) \simeq t(w')$, which contradicts to t being minimal. ■

Lemma 2.6 Given trees t, t' with $depth(t') < depth(t)$. Then $t' \not\preceq t$.

Proof Let $root(t) = v$, $depth(t) = n$ and $v \xrightarrow{r_1} \ldots \xrightarrow{r_n} v_{n+1}$ be a path of length n in t. Since t' is tree, i.e. an acyclic graph of depth $m < n$, there cannot be $w \in V(t')$ and a relation $\mathcal{S} \subseteq V(t') \times V(t)$ from t to t' such that (w, v_m) and (SC2) from Definition 2.2 holds. ■

Proposition 3.5 Let $\psi : T_{min} \to \mathcal{P}(T_{min})$ be a complete refinement operator. Then ψ is redundant.

PROOF We assume $\mathcal{K} = \emptyset$ and N_C contains A_1 and A_2. Since ψ is complete and its refinements are minimal, we have $t_\top \rightsquigarrow^* t_{A_1}$. Similarly, $t_\top \rightsquigarrow^* t_{A_1}$, $t_{A_1} \rightsquigarrow^* t_{A_1 \sqcap A_2}$, and $t_{A_2} \rightsquigarrow^* t_{A_1 \sqcap A_2}$.

We have $A_1 \not\sqsubseteq A_2$ and $A_2 \not\sqsubseteq A_1$, which means that $t_{A_1} \not\rightsquigarrow^* t_{A_2}$ and $t_{A_2} \not\rightsquigarrow^* t_{A_1}$.

Hence, $A_1 \sqcap A_2$ can be reached from $\top$ via a refinement chain going through A_1 and a different refinement chain not going through A_1, i.e. ψ is redundant. ■

A.2. Computing the maximal simulation

The proof of Lemma 2.3 gives rise to Algorithm 2. Given $\mathcal{EL}$ trees t_C and t_D, starting from the leaves of t_D, it computes bottom-up the maximal simulation from t_D to t_C in $\mathcal{O}(|t_C| \cdot |t_D|)$.

Algorithm 2 Computing the maximal simulation

Require: $\mathcal{EL}$ trees t_C, t_D
  $\mathcal{S} := \emptyset$
  for $i = 0$; $i \leq depth(t_D)$; $i := i + 1$ **do**
    for all $w \in V_D$ with $depth(w) = i$ **do**
      for all $v \in V_C$ **do**
        if (SC1) and (SC2) hold for (v,w) **then**
          $\mathcal{S} := \mathcal{S} \cup \{(v, w)\}$
        end if
      end for
    end for
  end for
  return$\mathcal{S}$

Lemma A.5 *Given $\mathcal{EL}$ trees t_C and t_D with roots x_C and x_D, Algorithm 2 computes the maximal simulation $\mathcal{S}$ from t_D to t_C.*

PROOF We prove the statement by induction on $n = depth(t_D)$. The induction base case follows obviously. For the induction step, let $\mathcal{S}_n$ be the relation obtained in the algorithm after iterating the outermost **for**-loop n times and let $\mathcal{S}$ be the relation obtained from the algorithm. It follows from the induction hypothesis that $\mathcal{S}_n$ is the maximal simulation from the subtree of every successor node of x_D to t_1. Hence for the maximal simulation $\mathcal{S}'$ from t_D to t_C, $\mathcal{S}' \setminus (V(t_C) \times \{x_D\}) \subseteq \mathcal{S}_n$. Now assume $(x_C, x_D) \in \mathcal{S}'$, but $(x_C, x_D) \notin \mathcal{S}$. Then for every r_2-successor y_D of x_D, there exist an r_1-successor y_C of x_C with $r_1 \sqsubseteq r_2$, and $(y_C, y_D) \in \mathcal{S}'$. However, $(y_C, y_D) \in \mathcal{S}_n \subseteq \mathcal{S}$ and consequently the pair (x_C, x_D) is also added in the last run of the outermost **for**-loop to $\mathcal{S}$. Hence, $\mathcal{S}$ is maximal. ■