Justification Masking in Ontologies

Matthew Horridge
Stanford University
Stanford, CA 94305, USA

Bijan Parsia and **Ulrike Sattler**
The University of Manchester
Oxford Road, Manchester
M13 9PL, United Kingdom

Abstract

This paper presents a characterisation of and definitions for the phenomenon of masking in the context of justifications for entailments in ontologies. In essence masking is present within a justification, over a set of justifications, or over a complete ontology when the number of justifications for an entailment does not reflect the number of reasons for that entailment. Four types of masking are defined in this paper: Internal Masking, Cross Masking, External Masking and Shared Core Masking. The results of an empirical study are presented which shows that the phenomenon of masking is prevalent throughout ontologies with non-trivial entailments in the NCBO BioPortal corpus. Out of 72 ontologies, 53 exhibited some form of masking, with 9 ontologies exhibiting internal masking, 23 ontologies exhibiting external masking, and 53 ontologies exhibiting shared core masking.

1 Introduction

Many open source and commercial ontology development tools such as Protégé, Swoop, The NeOn Toolkit and Top Braid Composer use justifications (Kalyanpur 2006) as a kind of explanation for entailments in ontologies. A *justification* for an entailment, also known as a *MinA* (Baader and Hollunder 1995; Baader, Peñaloza, and Suntisrivaraporn 2007), or a *MUPS* (Schlobach and Cornet 2003) if specific to explaining why a class name is unsatisfiable, is a minimal subset of an ontology that is sufficient for the given entailment to hold. More precisely, $\mathcal{J}$ is a justification for $\mathcal{O} \models \eta$ if $\mathcal{J} \subseteq \mathcal{O}$, $\mathcal{J} \models \eta$ and for all $\mathcal{J}' \subsetneq \mathcal{J}$ it is the case that $\mathcal{J}' \not\models \eta$. Justifications are a popular form of explanation in the OWL world and, as the widespread tooling support shows, have been used in preference to full blown proofs for explaining why an entailment follows from an ontology.

However, despite the popularity of justifications, they suffer from several problems stemming from their syntactic focus. Crucially, while all of the axioms in a justification are needed to support the entailment in question, there may be *parts* of these axioms that are not required for the entailment to hold (Horridge, Parsia, and Sattler 2008) . For example, consider $\mathcal{J} = \{A \sqsubseteq \exists R.B, Domain(R, C), C \sqsubseteq D \sqcap E\}$ which entails $A \sqsubseteq D$. While $\mathcal{J}$ is a justification for $A \sqsubseteq D$, and *all axioms* are required to support this entailment, there are *parts* of these axioms that are superfluous as far as the entailment is concerned: In the first axiom the filler of the existential restriction is superfluous, in the third axiom the conjunct E is superfluous for the entailment.

Given that justifications are not minimal with respect to their "parts" and thus with respect to their logical content, it is possible for the cardinality of the set of justifications to be different from the set of *reasons* for an entailment, that is, justifications can *mask* the set of reasons. For example, consider $\mathcal{J} = \{A \sqsubseteq \exists R.C \sqcap \forall R.C, D \equiv \exists R.C\}$ which entails $A \sqsubseteq D$. Clearly, $\mathcal{J}$ is a justification for $A \sqsubseteq D$. It is also noticeable that there are superfluous parts in this justification. Moreover, there are two distinct reasons why $\mathcal{J} \models A \sqsubseteq D$, the first being $\{A \sqsubseteq \exists R.C, \exists R.C \sqsubseteq D\}$ and the second being $\{A \sqsubseteq \exists R.\top \sqcap \forall R.C, \exists R.C \sqsubseteq D\}$. The work presented in the paper describes how masking can occur within a justification, over a set of justifications, and over a set of justifications plus axioms outside the set of justifications. The main problems identified with masking are (i) it can hamper understanding—not all reasons for an entailment may be salient to a person trying to understand the entailment, and (ii) it can hamper the design or choice of a repair plan—not all reasons for an entailment may be obvious, and if the plan consists of weakening and removing parts of axioms it may not actually result in a successful repair of the ontology in question.

2 Preliminaries

Description Logics This paper focuses on the Description Logic $\mathcal{SHOIQ}$. In general, Description Logics (DLs) are characterised as being decidable fragments of First Order Logic which have their own concise syntax. Rather than providing a token introduction to DLs the interested reader should look at a copy of the Description Logic Handbook and refer to (Horrocks and Sattler 2007) for specific details of $\mathcal{SHOIQ}$. In this paper, the symbols A and B are used to represent class names, C and D (possibly complex) $\mathcal{SHOIQ}$ concepts, R and S (possibly complex) $\mathcal{SHOIQ}$ roles and a and b individual names. A $\mathcal{SHOIQ}$ ontology consists of a set of axioms of the form $C \sqsubseteq D$, $R \sqsubseteq S$, $C(a)$ and $R(a, b)$. The axiom $C \equiv D$ is an abbreviation for $C \sqsubseteq D$ and $D \sqsubseteq C$. An ontology $\mathcal{O}$ (or an axiom α') entails an axiom α if every model of $\mathcal{O}$ (resp. α') is also a model of α. $\mathcal{O}$ entails α is written as $\mathcal{O} \models \alpha$ (resp. $\alpha' \models \alpha$). An axiom α is *weaker* than an axiom α' if both $\alpha' \models \alpha$ and $\alpha \not\models \alpha'$.

δ–The Structural Transformation One of the most important requirements in this work is the ability to identify the *occurrence* of a subconcept in some axiom—that is, the notion of a *subconcept with position*. This is achieved using a *structure preserving* transformation function δ, which removes the nesting of subconcepts by introducing fresh names, and produces axioms that are "small" and "flat". The transformation δ, which is presented below, is essentially the structural transformation which was originally described in Plaisted and Greenbaum (Plaisted and Greenbaum 1986).

In terms of Description Logics, the transformation takes a set of axioms $\mathcal{S}$, and produces a different set of axioms $\mathcal{S}' = \delta(\mathcal{S})$ that, while not equivalent to $\mathcal{S}$ is equi-consistent and equi-satisfiable to $\mathcal{S}$. That is, $\mathcal{S}'$ is consistent if and only if $\mathcal{S}$ is consistent, and a concept C is satisfiable with respect to $\mathcal{S}'$ if and only if it is satisfiable with respect to $\mathcal{S}$. Moreover, any model $\mathcal{I}$ of $\delta(\mathcal{S})$ is also a model of $\mathcal{S}$, and any model of $\mathcal{S}$ can be *extended* into a model of $\delta(\mathcal{O})$ by appropriate interpretation of the symbols that are in $\mathcal{S}'$ but not in $\mathcal{S}$.

In what follows a set of transformation rules that is used in the definition of δ is presented. The set of rules is *deterministic*. For a given ontology $\mathcal{O}$, the left hand side of each transformation rule matches an axiom in $\mathcal{O}$, and either replaces it with an axiom or a set of axioms that is defined by the right hand side of the rule.

Definition 1 *($\mathcal{SHOIQ}$ **Rewrite Rules for the Structural Transformation** δ)In the rewrite rules below, A_C, A_D and A_{C_i} are* fresh concept names *that are **not** in the signature of $\mathcal{O}$. The concepts C, D and C_i $(1 \leq i \leq n)$ are either complex concept expressions, concept names that **are** in the signature of $\mathcal{O}$, $\top$ or $\bot$*

T1	$C \equiv D \rightsquigarrow \{C \sqsubseteq D,\ D \sqsubseteq C\}$
R1	$S \equiv R \rightsquigarrow \{S \sqsubseteq R, R \sqsubseteq S\}$
A1	$C(a) \rightsquigarrow \{\{a\} \sqsubseteq C\}$
A2	$R(a,b) \rightsquigarrow \{\{a\} \sqsubseteq \exists R.\{b\}\}$
G1	$C \sqsubseteq D \rightsquigarrow \{A_C \sqsubseteq A_D,\ C \sqsubseteq A_C,\ A_D \sqsubseteq D\}$
P1	$A_D \sqsubseteq C_1 \sqcap \cdots \sqcap C_n \rightsquigarrow \{A_D \sqsubseteq C_i \mid 1 \leq i \leq n, C_i \neq \top\}$
P2	$A_D \sqsubseteq C_1 \sqcup \cdots \sqcup C_n \rightsquigarrow \{A_{C_i} \sqsubseteq C_i \mid 1 \leq i \leq n\} \cup \{A_D \sqsubseteq A_{C_1} \sqcup \cdots \sqcup A_{C_n}\}$
P3	$A_D \sqsubseteq \neg C, C \neq \bot \rightsquigarrow \{A_D \sqsubseteq \neg A_C,\ C \sqsubseteq A_C\}$
P4	$A_D \sqsubseteq \exists R.C, C \neq \top \rightsquigarrow \{A_D \sqsubseteq \exists R.A_C,\ A_C \sqsubseteq C\}$
P5	$A_D \sqsubseteq \forall R.C, C \neq \top \rightsquigarrow \{A_D \sqsubseteq \forall R.A_C,\ A_C \sqsubseteq C\}$
P6	$A_D \sqsubseteq \geq nR.C, C \neq \top \rightsquigarrow \{A_D \sqsubseteq \geq nR.A_C,\ A_C \sqsubseteq C\}$
P7	$A_D \sqsubseteq \leq nR.C, C \neq \bot \rightsquigarrow \{A_D \sqsubseteq \leq nR.A_C,\ C \sqsubseteq A_C\}$
N1	$C_1 \sqcap \cdots \sqcap C_n \sqsubseteq A_D \rightsquigarrow \{C_i \sqsubseteq A_{C_i} \mid 1 \leq i \leq n\} \cup \{A_{C_1} \sqcap \cdots \sqcap A_{C_n} \sqsubseteq A_D\}$
N2	$C_1 \sqcup \cdots \sqcup C_n \sqsubseteq A_D \rightsquigarrow \{C_i \sqsubseteq A_D \mid 1 \leq i \leq n, C_i \neq \bot\}$
N3	$\neg C \sqsubseteq A_D, C \neq \top \rightsquigarrow \{\neg A_C \sqsubseteq A_D,\ A_C \sqsubseteq C\}$
N4	$\exists R.C \sqsubseteq A_D, C \neq \bot \rightsquigarrow \{\exists R.A_C \sqsubseteq A_D,\ C \sqsubseteq A_C\}$
N5	$\forall R.C \sqsubseteq A_D, C \neq \bot \rightsquigarrow \{\forall R.A_C \sqsubseteq A_D,\ C \sqsubseteq A_C\}$
N6	$\geq nR.C \sqsubseteq A_D, C \neq \bot \rightsquigarrow \{\geq nR.A_C \sqsubseteq A_D,\ C \sqsubseteq A_C\}$
N7	$\leq nR.C \sqsubseteq A_D, C \neq \top \rightsquigarrow \{\leq nR.A_C \sqsubseteq A_D,\ A_C \sqsubseteq C\}$

Rule T1 rewrites concept equivalence axioms into general concept inclusion axioms. Rule R1 rewrites role equivalence axioms into two role inclusion axioms. For the sake of convenience, rules A1 and A2 rewrite ABox concept and role assertions into TBox axioms using nominals. Rules G1, P1 − P7 and N1 − N7 rewrite general concept inclusion axioms into multiple axioms, flattening out all nested concept expressions. They do this by introducing fresh names for subconcepts and introducing "defining axioms" for these fresh names.

Definition 2 ***(The Structural Transformation** δ **for** $\mathcal{SHOIQ}$ **)** For a given set of $\mathcal{SHOIQ}$ axioms $\mathcal{O}$, $\mathcal{O}' = \delta(\mathcal{O})$ is the result of* exhaustively *applying the rewrite rules given in Definition 1.*

δ-Transformation Axiom Forms Applying the structural transformation δ to a set of axioms $\mathcal{S}$ results in a set of small flat axioms $\delta(\mathcal{S})$. Each axiom in $\delta(\mathcal{S})$ is of one of the forms shown in Definition 3, where the symbols N_*, P_*, and A represent concept names, R a $\mathcal{SHOIQ}$ role, o an individual name, and n a positive integer. Each concept name P_c represents a *positive occurrence* of some subconcept C in the original set of axioms, with an axiom of the form $P_c \sqsubseteq C$ "defining" that occurrence. Each concept name N_c represents a *negative occurrence* of some subconcept C in the original set of axioms, with each axiom of the form $C \sqsubseteq N_c$ "defining" that occurrence.

Definition 3 *($\mathcal{SHOIQ}$ **δ-transformation axiom forms**) For a set $\mathcal{S}$ of $\mathcal{SHOIQ}$ axioms, each axiom $\alpha \in \delta(\mathcal{S})$ must be one of the forms, where the symbols N_*, P_*, and A represent concept names, R a $\mathcal{SHOIQ}$ role, o an individual name, and n a positive integer:*

A1	$N_1 \sqsubseteq P_1$
P1	$P_c \sqsubseteq P_1 \sqcup \cdots \sqcup P_n$
P2	$P_c \sqsubseteq \neg N_1$
P3	$P_c \sqsubseteq \{o\}$
P4	$P_c \sqsubseteq A$
P5	$P_c \sqsubseteq \exists R.P_1$
P6	$P_c \sqsubseteq \forall R.P_1$
P7	$P_c \sqsubseteq \geq nR.P_1$
P8	$P_c \sqsubseteq \leq nR.N_1$
N1	$N_1 \sqcap \cdots \sqcap N_n \sqsubseteq N_c$
N2	$\neg N_1 \sqsubseteq N_c$
N3	$\{o\} \sqsubseteq N_c$
N4	$A \sqsubseteq N_c$
N5	$\exists R.N_1 \sqsubseteq N_c$
N6	$\forall R.N_1 \sqsubseteq N_c$
N7	$\geq nR.N_1 \sqsubseteq N_c$
N8	$\leq nR.P_1 \sqsubseteq N_c$

O1	trans(R)

$\mathcal{SHOIQ}$ **Syntactic Isomorphism** Intuitively, an axiom α' is syntactically isomorphic to another axiom α'' if there is an injective renaming ρ of the signature, with $\top$ and $\bot$, of α' so that $\rho(\alpha')$ is structurally equal to α''. For example, $\alpha' = A \sqsubseteq \exists R.B$ is isomorphic to $\alpha'' = F \sqsubseteq \exists S.B$, since A and R in α' can be renamed to F and S respectively, to make α' structurally equal to α''. Definition 4 captures what it means to be isomorphic for an axiom that occurs in the set of axioms $\delta(\mathcal{S})$ where $\mathcal{S}$ is a set of $\mathcal{SHOIQ}$ axioms. Note that structural equality is used as the ordering of conjuncts in a conjunction and disjuncts in a disjunction is unimportant here.

Definition 4 ($\mathcal{SHOIQ}$ δ-Isomorphism) *Two $\mathcal{SHOIQ}$ axioms, α' and α'' are δ-isomorphic if α' and α'' are both of **one** of the forms given in Definition 3, and there is a injective renaming of each $N' \in$ signature$(\alpha') \cup \{\top, \bot\}$ into a name $N'' \in$ signature$(\alpha'') \cup \{\top, \bot\}$.*

Laconic Justifications Laconic justifications centre around the notions of superfluity and weakness. Roughly speaking, a justification $\mathcal{J}$ for an entailment η is laconic if: (1) $\mathcal{J}$ does not contain any axioms that contain any sub-concept *occurrences* (i.e. sub-concepts at specific positions) that are superfluous for η, and (2) $\mathcal{J}$ does not contain any sub-concept occurrences that could be weakened while preserving η. More precisely (where $\mathcal{O}^\star$ is the deductive closure of $\mathcal{O}$),

Definition 5 ***(Laconic Justification)*** *Let $\mathcal{O}$ be an ontology such that $\mathcal{O} \models \eta$, $\mathcal{J}$ is a* laconic justification *for η over $\mathcal{O}$ if:*

1. *$\mathcal{J}$ is a justification for η in $\mathcal{O}^\star$*
2. *For every $\mathcal{J}_\delta \subsetneq \delta(\mathcal{J})$ it is the case that $\mathcal{J}_\delta \not\models \eta$*
3. *For each $\alpha \in \delta(\mathcal{J})$ there is no α' such that*
   (a) $\alpha \models \alpha'$ and $\alpha' \not\models \alpha$ (α' is weaker than α)
   (b) α' is δ-isomorphic to α
   (c) $\delta(\mathcal{J}) \setminus \{\alpha\} \cup \{\alpha'\}$ is a justification for η in $(\delta(\mathcal{O}))^\star$

Non-Trivial Entailments Finally, the empirical study, which is described later in the paper, uses the notion of *non-trivial entailments*, which are defined as follows:

Definition 6 (Non-Trivial Entailments) *Given an ontology $\mathcal{O}$ and entailment η such that $\mathcal{O} \models \eta$, η is a* non-trivial *entailment if $\mathcal{O} \setminus \{\eta\} \models \eta$.*

Intuitively, an entailment η is non-trivial if it is either not asserted in $\mathcal{O}$ (i.e. $\eta \notin \mathcal{O}$), or, there are multiple justifications for $\mathcal{O} \models \eta$ (one of which is $\{\eta\}$).

3 Intuitions Behind Masking

This paper characterises and defines four important types of masking: *Internal Masking*, *Cross Masking*, *External Masking* and *Shared Core Masking*. The intuitions behind these types of masking are explained below.

Internal Masking Internal masking refers to the phenomena where there are multiple reasons within a single justification as to why the entailment in question holds. An example of internal masking is shown below.

$$\mathcal{O} = \{A \sqsubseteq B \sqcap \neg B \sqcap C \sqcap \neg C\} \models A \sqsubseteq \bot$$

There is a single regular justification for $\mathcal{O} \models A \sqsubseteq \bot$, namely $\mathcal{O}$ itself. However, within this justification there are, intuitively, two reasons as to why $\mathcal{O} \models A \sqsubseteq \bot$, the first being $\{A \sqsubseteq B \sqcap \neg B\}$ and the second being $\{A \sqsubseteq C \sqcap \neg C\}$.

Cross Masking Intuitively, cross masking is present within a set of justifications for an entailment when parts of axioms from one justification combine with parts of axioms from another justification in the set to produce new reasons for the given entailment. For example, consider the following ontology.

$$\begin{aligned}\mathcal{O} = \{&A \sqsubseteq B \sqcap \neg B \sqcap C\\ &A \sqsubseteq D \sqcap \neg D \sqcap \neg C\} \models A \sqsubseteq \bot\end{aligned}$$

There are two justifications for $\mathcal{O} \models A \sqsubseteq \bot$, namely $\mathcal{J}_1 = \{A \sqsubseteq B \sqcap \neg B \sqcap C\}$ and $\mathcal{J}_2 = \{A \sqsubseteq D \sqcap \neg D \sqcap \neg C\}$. However, part of the axiom in $\mathcal{J}_1$, namely $A \sqsubseteq C$ may combine with part of the axiom in $\mathcal{J}_2$, namely $A \sqsubseteq \neg C$ to produce a further reason: $\mathcal{J}_3 = \{A \sqsubseteq C, A \sqsubseteq \neg C\}$.

External Masking While internal masking and cross masking take place over a set of "regular" justifications for an entailment, external masking involves parts of axioms from a regular justification combining with parts of axioms from an ontology (intuitively the axioms outside of the set of regular justifications) to produce further reasons for the entailment in question. Consider the example below,

$$\begin{aligned}\mathcal{O} = \{&A \sqsubseteq B \sqcap \neg B \sqcap C\\ &A \sqsubseteq \neg C\} \models A \sqsubseteq \bot\end{aligned}$$

There is just one justification for $\mathcal{O} \models A \sqsubseteq \bot$, however, although $A \sqsubseteq \neg C$ intuitively plays a part in the unsatisfiability of A it will never appear in a justification for $\mathcal{O} \models A \sqsubseteq \bot$. When $\mathcal{O}$ is taken into consideration, there are two salient reasons for $A \sqsubseteq \bot$, the first being $\{A \sqsubseteq B \sqcap \neg B\}$ and the second being $\{A \sqsubseteq C, A \sqsubseteq \neg C\}$

Shared Core Masking Finally, two justifications share a *core* if after stripping away the superfluous parts of axioms in each justification the justifications are essentially structurally equal. Consider the example below,

$$\begin{aligned}\mathcal{O} = \{&A \sqsubseteq B \sqcap \neg B \sqcap C\\ &A \sqsubseteq B \sqcap \neg B\} \models A \sqsubseteq \bot\end{aligned}$$

There are two justifications for $\mathcal{O} \models \eta$, $\mathcal{J}_1 = \{A \sqsubseteq B \sqcap \neg B \sqcap C\}$ and $\mathcal{J}_2 = \{A \sqsubseteq B \sqcap \neg B\}$. However, $\mathcal{J}_1$ can be reduced to the laconic justification $\{A \sqsubseteq B \sqcap \neg B\}$ (since C is irrelevant for the entailment), which is structurally equal to $\mathcal{J}_2$. With regular justifications, it appears that there are more reasons for the entailment, when in fact each justification is precisely the same reason.

Masking Due to Weakening The above intuitions have been illustrated using simple propositional examples. However, it is important to realise that masking is not just concerned with Boolean parts of axioms. *Weakest parts* of axioms must also be taken into consideration. For example, consider

$$\begin{aligned}\mathcal{O} = \{ & A \sqsubseteq \geq 2R.C \\ & A \sqsubseteq \geq 1R.D \\ & C \sqsubseteq \neg D\} \models A \sqsubseteq \geq 2R\end{aligned}$$

There is one regular justification for $\mathcal{O} \models A \sqsubseteq \geq 2R$ namely, $\mathcal{J}_1 = \{A \sqsubseteq \geq 2R.C\}$. However, there are intuitively two reasons for this entailment. The first is described by the justification obtained as a weakening of $\mathcal{J}_1$, and is $\mathcal{J}_2 = \{A \sqsubseteq \geq 2R\}$. The second is obtained by weakening the first axiom in $\mathcal{O}$ and combining it with the second and third axioms in $\mathcal{O}$ to give $\{A \sqsubseteq \geq 1R.C, A \sqsubseteq \geq 1R.D, C \sqsubseteq \neg D\}$. Of course, any form of masking can be due to weakening.

4 Masking Defined

With the above intuitions and desiderata in mind the notion of masking can be made more concrete. In the spirit of laconic justifications, the basic idea is to pull apart the axioms in a justification, set of justifications and an ontology, compute constrained weakenings of these parts (inline with the definition of laconic justifications), and then to check for the presence and number of laconic justifications within the set of regular justifications for an entailment with respect to these parts and their weakenings.

Parts and Their Weakenings First, it is necessary define a function $\delta^+(\mathcal{S})$, which maps a set of axioms $\mathcal{S}$ to a set of axioms composed from the union of $\delta(\mathcal{S})$ with the constrained weakenings of axioms in $\delta(\mathcal{S})$. The weakenings of axioms are constrained in accordance with Definition 5(3). For an axiom $\alpha \in \delta(\mathcal{S})$, a weakening α' of α is contained in $\delta^+(\mathcal{S})$ only if α' is of the same form as α—i.e. α' is δ-isomorphic to α.

Definition 7 **(δ^+)** *For a set of $\mathcal{SHOIQ}$ axioms, $\mathcal{S}$,*

$$\begin{aligned}\delta^+(\mathcal{S}) := \delta(\mathcal{S}) \cup \{\alpha \mid \exists \alpha' \in \delta(\mathcal{S}) \textit{ s.t. } & \\ & \alpha' \models \alpha \textit{ and} \\ & \alpha \not\models \alpha' \textit{ and} \\ & \alpha' \textit{ is } \delta\textit{-isomorphic to } \alpha\}\end{aligned}$$

Lemma 1 (δ^+justificatory finiteness) *For a finite set of axioms $\mathcal{S}$, the set of justifications for an entailment in $\delta^+(\mathcal{S})$ is finite.*

Next, a function which filters out laconic justifications for an entailment from a set of all justifications for the entailment is defined in Definition 8.

Definition 8 (Laconic Filtering) *For a set of axioms $\mathcal{S} \models \eta$, $\mathsf{laconic}(\mathcal{S}, \eta)$ is the set of justifications for $\mathcal{S} \models \eta$ that are laconic over $\mathcal{S}$.*

It should be noted that because of Lemma 1, the set of justifications $\mathsf{laconic}(\delta^+(\mathcal{S}), \eta)$ is finite.

Masking Definitions With the definition of δ^+ and the definition of laconic filtering in hand, the various types of masking can now be defined.

Definition 9 (Internal Masking) *For a justification $\mathcal{J}$ for $\mathcal{O} \models \eta$, internal masking is present within $\mathcal{J}$ if*

$$\left|\mathsf{laconic}(\delta^+(\mathcal{J}), \eta)\right| > 1$$

Theorem 1 *Internal masking is not present within a laconic justification.*

Let $\mathcal{O} \models \eta$ and $\mathcal{J}_1, \ldots, \mathcal{J}_n$ be the set of all justifications for $\mathcal{O} \models \eta$. Cross masking and External masking are then defined as follows:

Definition 10 (Cross Masking) *For two justifications $\mathcal{J}_i$ and $\mathcal{J}_j$, cross masking is present within $\mathcal{J}_i$ and $\mathcal{J}_j$ if*

$$\left|\mathsf{laconic}\big(\delta^+(\mathcal{J}_i \cup \mathcal{J}_j), \eta\big)\right| > \Big(\left|\mathsf{laconic}\big(\delta^+(\mathcal{J}_i), \eta\big)\right| + \left|\mathsf{laconic}\big(\delta^+(\mathcal{J}_j), \eta\big)\right|\Big)$$

Definition 11 (External Masking) *External masking is present if*

$$\left|\mathsf{laconic}(\delta^+(\mathcal{O}), \eta)\right| > \left|\mathsf{laconic}(\delta^+(\bigcup_{i=1}^{i=n} \mathcal{J}_i), \eta)\right|$$

Definition 12 (Shared Cores) *Two justifications $\mathcal{J}_i$ and $\mathcal{J}_j$ for $\mathcal{O} \models \eta$, share a core if there is a justification $\mathcal{J}'_i \in \mathsf{laconic}(\delta^+(\mathcal{J}_i), \eta)$ and a justification $\mathcal{J}'_j \in \mathsf{laconic}(\delta^+(\mathcal{J}_j), \eta)$ and an injective renaming ρ of terms not in $\mathcal{O}$ such that $\rho(\mathcal{J}'_i) = \mathcal{J}'_j$.*

5 Masking in the Field

The first part of this paper has focused on pinning down the various notions of masking and making these notions more precise in terms of definitions. The second part details a study examining the prevalence of masking in real ontologies. Ontologies for the study were obtained from the NCBO BioPortal. Due to limited space, algorithms and optimisations for examining masking, corpus description/makeup, and experimental procedure are not presented here. A detailed presentation and discussion may be found in (Horridge 2011). The BioPortal ontology repository was accessed on the 12th March 2011. A total of 218 OWL compatible (OWL and OBO) ontology documents were processed. For each ontology, non-trivial entailments that were direct subsumptions between concept names, and direct class assertions between concept names and individual names were computed. There were 72 ontologies that contained non-trivial entailments of this form. For each entailment, regular and masked laconic justifications were computed and the numbers compared.

Figure 1 shows the mean number of regular and laconic justifications per entailment per ontology. Figure 2 shows a bubble plot that depicts entailments where the number of regular justifications does not equal the number of masked justifications, and so gives some idea of masking over the whole BioPortal corpus. The x-axis shows the number of regular justifications and the y-axis shows the number of

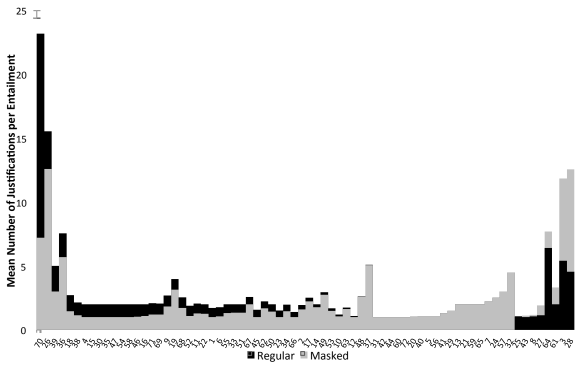

Figure 1: Mean number of regular and masked justifications per entailment

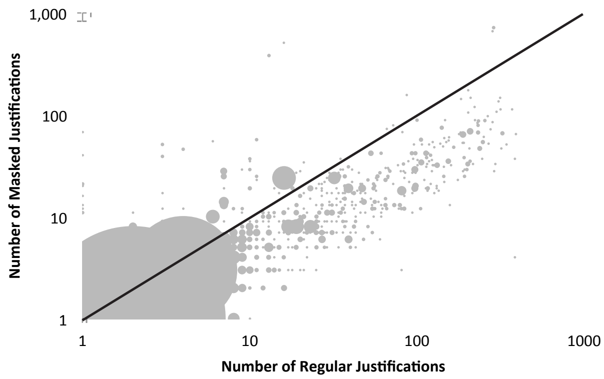

Figure 2: The Effect of Masking – Ratios of regular justifications to masked justifications

masked justifications. The size of the bubbles reflects the number of entailments with a particular ratio of regular to masked justifications. For example, the large bubble in the lower left corner represents entailments which have two regular justifications but only one masked justification, of which there are 5,447 entailments.

As can be seen from Figure 1 and Figure 2 the phenomena of masking, be it internal, external or shared core masking is prevalent throughout the ontologies in the BioPortal corpus. Most of the ontologies (53 in total out of 72) exhibit some kind of masking. Specifically, there are 9 ontologies that exhibit internal masking, 23 ontologies that exhibit external masking, and 53 ontologies that exhibit shared core masking.

The bubble plot in Figure 2 shows where the vast majority of masking cases lie[1]—shared core masking dominates, with plenty of entailments that have between 2 and 10 regular justifications but between only 1 and 5 masked laconic justifications. However, there are also plenty of examples of external masking, with some of them being quite extreme. For example, the small bubbles that occur in the upper middle of Figure 2 represent entailments that have around 14-18 regular justifications but around 400 and 500 hundred masked laconic justifications.

[1]It should be noted that Figure 2 is a log-log plot and any bubbles which lie at a distance from the diagonal centre line represent fairly large differences.

6 Conclusions

This paper has presented and defined justification masking. In essence, masking is a phenomenon that results in the number of justifications for an entailment not reflecting the *number of reasons* for an entailment. Four specific types of masking have been defined in terms of the cardinality of underlying reasons for the target entailment: Internal masking, Cross-masking, External masking, and Shared-Core masking. There is strong evidence to suggest that each type of masking occurs for non-trivial entailments in real world ontologies. Each has significant implications for understanding entailments whether from a direct user perspective (e.g., debugging rogue entailments) or indirect (e.g., presenting metrics about amount of logical content). External masking and Shared-Core masking are especially striking. External masking shows that justifications do not always capture all axioms which are relevant to an entailment. Shared-Core masking shows that presenting justifications to users can results in significant wasted effort as the user is presented with prima facie more reasons than there in fact are.

References

Baader, F., and Hollunder, B. 1995. Embedding defaults into terminological knowledge representation formalisms. *Journal of Automated Reasoning* 14(1):149–180.

Baader, F.; Peñaloza, R.; and Suntisrivaraporn, B. 2007. Pinpointing in the description logic $\mathcal{EL}$. In *KI 2007: Advances in Artificial Intelligence, Osnabrück, Germany*, volume 4667 of *LNCS*, 52–67. Springer.

Horridge, M.; Parsia, B.; and Sattler, U. 2008. Laconic and precise justifications in OWL. In *ISWC 2008, 7th International Semantic Web Conference, Karlsruhe, Germany, October 26-30, 2008.*, volume 5318 of *LNCS*, 323–338. Springer.

Horridge, M. 2011. *Justification Based Explanation in Ontologies*. Ph.D. Dissertation, School of Computer Science, The University of Manchester.

Horrocks, I., and Sattler, U. 2007. A tableau decision procedure for $\mathcal{SHOIQ}$. *Journal of Automated Reasoning* 39(3):249–276.

Kalyanpur, A. 2006. *Debugging and Repair of OWL Ontologies*. Ph.D. Dissertation, The Graduate School of the University of Maryland.

Plaisted, D. A., and Greenbaum, S. 1986. A structure-preserving clause form translation. *Journal of Symbolic Computation* 2(3):293–304.

Schlobach, S., and Cornet, R. 2003. Non-standard reasoning services for the debugging of description logic terminologies. In Gottlob, G., and Walsh, T., eds., *IJCAI-03, Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence, Acapulco, Mexico, August 9-15, 2003*, 355–362. Morgan Kaufmann.