Inspecting Regularities in Ontology Design Using Clustering

Eleni Mikroyannidi, Luigi Iannone, Robert Stevens, and Alan Rector

The University of Manchester
Oxford Road, Manchester, M13 9PL
{mikroyannidi,iannone,stevens,rector}@cs.manchester.ac.uk

Abstract. We propose a novel application of clustering analysis to identify regularities in the usage of entities in axioms within an ontology. We argue that such regularities will be able to help to identify parts of the schemas and guidelines upon which ontologies are often built, especially in the absence of explicit documentation. Such analysis can also isolate irregular entities, thus highlighting possible deviations from the initial design. The clusters we obtain can be fully described in terms of generalised axioms that offer a synthetic representation of the detected regularity. In this paper we discuss the results of the application of our analysis to different ontologies and we discuss the potential advantages of incorporating it into future authoring tools.

1 Introduction

Ontologies are often built according to guidelines or schemas that give rise to repeating regularities in the use of entities in axioms. Recognising those regularities is important in understanding the ontology and assuring that it conforms to the schemas. A regular ontology shows organisation of the knowledge and thus its coherence. For example, in the wine ontology¹, all the wines at the individual level are described in similar ways; they have a particular type of wine and there are property assertions defining their body, origin, maker, flavour etc. The inspection of these regularities can give an insight into the construction of the ontology. By looking at the description of some wines, the user can have an insight about the template that was used to describe them. Deviations from this regular design might either be legitimate exceptions in the regularities or defects in modelling. For example, the individual TaylorPort does not have any property assertions referring to its origin, as this is defined on its type (Port SubClassOf locatedIn value PortugalRegion). Having a mechanism that can help to isolate such deviations could be a useful tool for quality assurance.

It is, however, difficult to trace irregularities and regularities by eye, especially in big ontologies. The reasoner can check the satisfiability of concepts in the ontology, and there are tools [9,5] that provide explanations for unsatisfiabilities

¹ http://www.w3.org/TR/owl-guide/wine

L. Aroyo et al. (Eds.): ISWC 2011, Part I, LNCS 7031, pp. 438–453, 2011.

[©] Springer-Verlag Berlin Heidelberg 2011

and other entailments. However, they cannot trace irregularities in the design as these are not logical errors.

Ontology environments such as Protégé-4, Swoop, NeOn Toolkit and Top Braid Composer provide some structural information about the ontology through visualisations, such as hierarchies and panels showing the usage of the entities. The task of spotting regularities in the usage of axioms in entity descriptions is, however, not supported beyond this basic exposure of usage.

Ontology engineering has adopted the idea of design patterns [4] to capture accepted modelling solutions to common issues [2,3]. Patterns of axioms, however, can exist throughout an ontology without being an accepted design pattern. In this paper we focus on the general notion of patterns. In the remainder of the paper we will refer to them as *regularities*. Such a regularity is illustrated in the following axioms in the wine ontology:

```
lpha_1 = {\sf PinotBlanc}\ EquivalentTo\ {\sf wine\ and}
({\sf madeFromGrape\ value\ PinotBlancGrape})\ {\sf and\ (madeFromGrape\ max\ 1\ Thing)}
lpha_2 = {\sf CabernetSauvignon}\ EquivalentTo\ {\sf wine\ and\ (madeFromGrape\ value\ CabernetSauvignonGrape)\ and\ (madeFromGrape\ max\ 1\ Thing)}
```

The regularity can be expressed with the abstract axiom:

```
\alpha = \text{?xWine } \textit{EquivalentTo} \text{ wine and}
(\mathsf{madeFromGrape} \ \mathbf{some} \ \text{?xGrape}) \ \mathbf{and} \ (\mathsf{madeFromGrape} \ \mathbf{max} \ 1 \ \mathsf{Thing})
```

 $\label{eq:where and action} where ?xWine = \{PinotBlanc, CabernetSauvignon \}, ?xGrape = \{PinotBlancGrape, CabernetSauvignonGrape\} \ are \ variables \ holding \ similar \ entities.$

We present a novel framework for inspecting such regularities and clustering entities in the ontology according to these kinds of regularities. Clustering is a common scientific approach for identifying similarities [13]. The proposed method is based on capturing similar usage of the entities; such similar entities are expected to result in the same cluster. The description of the clusters is provided in abstract forms; axioms containing meaningful placeholders (e.g. ?xWine EquivalentTo wine and (madeFromGrape some ?xGrape) and (madeFromGrape max 1 Thing)). With such a framework, tools could help the inspection of regularities when attempting to comprehend ontologies, where information on the design style and schemas is often poorly documented and not represented explicitly in the ontology itself.

Efforts for developing or transforming ontologies using patterns are reported in [12], [2]. In [7], the use of the Ontology Pre-Processor Language (OPPL) as a means of embedding knowledge patterns in OWL ontologies is presented. However, little work is known in the field of regularity detection in ontologies. In [11] an analysis of collections of OWL ontologies with the aim of determining the frequency of several combined name and graph patterns is described. However, this work mainly focuses on lexical patterns, with some additional examination on how these apply in the taxonomy hierarchy. It does not explicitly involve identification of regularities in axiom usage. There appears to be no framework

that can identify regularities in the axioms of an ontology and highlight entities that have been designed according to these regularities.

2 Clustering Framework

The purpose of cluster analysis is to divide data into groups (clusters) that are meaningful, useful or both [13]. Our particular problem is to partition the set of entities in an ontology according to their usage, i.e.: entities in the same cluster occur with similar axioms playing similar roles. Cluster analysis relies on the notion of distance to quantify how similar (or dissimilar) and, therefore, how close or far apart two entities are in the clustering space. In the literature there are several clustering techniques; we chose agglomerative clustering and we report an informal description of the algorithm we used in Algorithm 1. In our algorithm, the implementation of step 1, and in particular, the distance adopted to compute the proximity matrix is the most important aspect of the implementation and the main focus of the rest of this section; steps 2–7 are typical steps for agglomerative cluster analysis and are not explained further. However, in our implementation, the algorithm will continue agglomerating until the distance between all possible pairs of elements in the two closest clusters is less than 1.

Algorithm 1. Clustering algorithm

Require: A set of entities. Ensure: A set of clusters.

- 1: Compute the proximity matrix {matrix containing the values of the distance between all entities}
- 2: Assign to the current set of clusters the set of singleton clusters, each representing an input entity.
- 3: repeat
- 4: Merge the closest two clusters, replace them with the result of the merge in the current set of clusters.
- 5: Update the proximity matrix, with the new distance between the newly created cluster and the original ones.
- 6: **until** The stopping criterion is met
- 7: return The current set of clusters.

In the approach taken here, the calculation of the distance is based on the similarity of the structure of the axioms in the ontology. For axiom comparison, we transform them into a more abstract form following a *place-holder replacement policy*. We will define the notion of similar structure of axioms more formally in the following, and show how this leads to the particular distance we adopted for our cluster analysis.

Comparing axioms: Let us consider the Pizza Ontology². Its main scope is pizzas and their toppings along with other information such as the country of

 $^{^2}$ http://www.co-ode.org/ontologies/pizza/

origin for each pizza, the spiciness of some toppings, or their classification as vegetables, fish, or meat. Among other things, in this ontology, we observe that all the topping classes are used as fillers in axioms like:

```
aPizza SubClassOf hasTopping some aTopping or anotherTopping or ...)
```

In other words, classes like MozzarellaTopping and TomatoTopping seem similar because they appear as fillers of an existential restriction on the property hasTopping within a sub-class axiom where the left-hand side is a pizza. They also appear as disjuncts in a universal restriction on the same property in another sub-class axiom whose left-hand side is, again, a pizza. Likewise, pizzas in this ontology tend to appear on the left-hand side of sub-class axioms describing their toppings, base, and country of origin. Therefore, our cluster analysis should, in the case of the pizza ontology, put together all toppings in a single cluster, pizzas in another, and countries of origin in a third one and so on. More formally, we need to introduce a distance that quantifies the difference between the usage of two entities in a set of axioms (ontology).

Definition 1. [Place-holder replacement] Let \mathcal{O} be an ontology and let $\Phi = \{ \text{owlClass}, \text{owlObjectProperty}, \text{owlDataProperty}, \text{owlAnnotationProperty}, \text{owlIndividual}, * \} be a set of six symbols that do not appear in the signature}^3 of <math>\mathcal{O} - \text{sig}(\mathcal{O})$. A place-holder replacement is a function $\phi : \text{sig}(\mathcal{O}) \to \text{sig}(\mathcal{O}) \cup \Phi$ satisfying the following constraints: Consider $e \in \mathcal{O}$ then $\phi(e) = e^{-i\theta}$

- $e \ or * or \ owlClass \ if \ e \ is \ a \ class \ name;$
- e or * or owlObjectProperty if e is a object property name;
- − e or * or owlDataProperty if e is a data property name;
- *e or* * *or owl*AnnotationProperty *if e is a annotation property name;*
- − e or * or owlIndividual if e is an individual property name.

We define the particular placeholder replacement ϕ^S as $\phi^S(e)$ =

- owlClass if e is a class name;
- $\ \textit{owlObjectProperty} \ \textit{if} \ e \ \textit{is} \ \textit{an object property name};$
- owlDataProperty if e is a data property name;
- $\ \textit{owl} \textit{AnnotationProperty} \ \textit{if} \ e \ \textit{is} \ \textit{an annotation property name};$
- owlIndividual if e is an individual property name.

Definition 2. [Place-holder replacement in axioms] Let \mathcal{O} be an ontology, $\alpha \in \mathcal{O}$ one of its axioms and ϕ a place-holder replacement function. We define ϕ_{Ax} as a function that for the input α returns a new axiom by applying ϕ to all the entities $e \in sig(\alpha)$.

Example 1. Let our \mathcal{O} be the Pizza ontology mentioned above and let us define ϕ as follows. $\forall e \in \mathcal{O}, \ \phi(e) =$

³ For signature here we mean the set of class names, data/object/annotation property names, individuals referenced in the axioms of an ontology \mathcal{O} .

- * if $e \in \{\text{Margherita}, \text{Capricciosa}\}^4;$ - $\phi^S(e)$ otherwise;

Let us now compute the values of $\phi_{Ax}(\alpha)$ for some of the axioms in \mathcal{O}

- $-\alpha = Margherita \ DisjointWith \ Cajun, \ \phi_{Ax}(\alpha) = *DisjointWith \ owlClass;$
- $-\alpha = \mathsf{Capricciosa}\ \mathsf{DisjointWith}\ \mathsf{Cajun},\ \phi_{Ax}(\alpha) = *\mathsf{DisjointWith}\ \mathsf{owlClass};$
- $-\alpha = \text{Margherita } SubClassOf \text{ hasTopping some TomatoTopping,}$ $\phi_{Ax}(\alpha) = *SubClassOf \text{ owlObjectProperty some owlClass;}$
- $-\alpha = \text{Capricciosa } \textit{SubClassOf} \text{ hasTopping some TomatoTopping,}$ $\phi_{Ax}(\alpha) = *\textit{SubClassOf} \text{ owlObjectProperty some owlClass;}$

We have defined the replacement function that transforms the axioms into abstractions and we can proceed with the measure of the distance.

Distance measure: We define the distance function as follows:

Definition 3. [Distance] Let \mathcal{O} be an ontology, e_1 and e_2 be two entities from $sig(\mathcal{O})$ and ϕ a place-holder replacement function. We denote $Ax_{\phi}(e)$ the set $\{\phi_{Ax}(\alpha), \alpha \in \mathcal{O}, e \in sig(\alpha)\}$, i.e. the set of pace-holder replacements for the axioms in \mathcal{O} that reference e.

We define the distance between the two entities, $d_{\phi}(e_1, e_2)$ as:

$$d(e_1, e_2) = \frac{|(Ax_{\phi}(e_1) \cup Ax_{\phi}(e_2)| - |Ax_{\phi}(e_1) \cap Ax_{\phi}(e_2)|}{|(Ax_{\phi}(e_1) \cup Ax_{\phi}(e_2)|}$$

From this we can observe that $\forall \mathcal{O}, \phi, e_1, e_2 : 0 \leq d_{\phi}(e_1, e_2) \leq 1$. The place-holder replacement function ϕ is a way to control the granularity of our distance.

Example 2. Let our \mathcal{O} be the Pizza ontology again and let us define ϕ_1 as follows. $\forall e \in \mathcal{O}, \phi_1(e) =$

- $* if e \in {TomatoTopping, PizzaBase};$
- $-\phi^S(e)$ otherwise;

Let us now compute the values of $\phi_{1Ax}(\alpha)$ for a pair of axioms in \mathcal{O}

- $-\alpha = Margherita$ *SubClassOf* hasTopping **some** TomatoTopping,
 - $\phi_{1Ax}(\alpha) = \text{owlClass } SubClassOf \text{ owlObjectProperty some *};$

 $\phi_{1Ax}(\alpha) = \text{owlClass} \ \textit{SubClassOf} \ \text{owlObjectProperty some *}.$

This means that $d_{\phi 1}(\mathsf{TomatoTopping}, \mathsf{PizzaBase}) < 1 \text{ as } |\mathsf{Ax}_{\phi_1}(e_1) \cap \mathsf{Ax}_{\phi_1}(e_2)| > 0$ and, therefore, $|(\mathsf{Ax}_{\phi_1}(e_1) \cup \mathsf{Ax}_{\phi_1}(e_2)| - |\mathsf{Ax}_{\phi_1}(e_1) \cap \mathsf{Ax}_{\phi_1}(e_2)| < |(\mathsf{Ax}_{\phi_1}(e_1) \cup \mathsf{Ax}_{\phi_1}(e_2)|.$

The consequence would be that our distance $d_{\phi 1}$ does not separate as cleanly as possible TomatoTopping (and likewise several sub-classes of PizzaTopping) from PizzaBase. Let us now compare it with another place-holder replacement function, ϕ_2 , defined as follows:

⁴ The * placeholder represents entities, which distance is computed. We denote this placeholder for keeping track of the position of the entities in the referencing axioms.

```
-* if e \in \{\mathsf{TomatoTopping}, \mathsf{PizzaBase}\};
```

-e if e is a object property name;

 $-\phi^{S}(e)$ otherwise;

Then our ϕ_{2Ax} for the same values for α will be:

- $-\alpha = \text{Margherita } SubClassOf \text{ hasTopping some TomatoTopping,}$ $\phi_{2Ax}(\alpha) = \text{owlClass } SubClassOf \text{ hasTopping some *;}$
- $-\alpha =$ Pizza SubClassOf hasBase **some** PizzaBase, $\phi_{2Ax}(\alpha) =$ owlClass SubClassOf hasBase **some ***

This will keep $d_{\phi 2}(\mathsf{TomatoTopping}, \mathsf{PizzaBase}) = 1$

Changing the granularity of the place-holder replacement function produces more or less sensitive distance functions. The two extremes are replacing every entity with a place-holder or not replacing any of them. Whilst the former produces a distance that is far too tolerant and puts together entities that seem unrelated, the latter will most likely result in a distance that scores 1 (maximal distance) for most entity pairs. In this work we propose a tradeoff where we delegate the decision of whether to replace an entity in an axiom to a measure of its **popularity** with respect to the other entities in the same kind of axiom within the ontology. More formally:

Definition 4 (Popularity). Let \mathcal{O} be an ontology, $e \in sig(\mathcal{O})$ an entity. The place-holder replacement function ϕ^{S}_{Ax} for the axioms of \mathcal{O} will extract the structure of each axiom.

Given an axiom $\alpha \in \mathcal{O}$, let us define the set $\mathsf{Ax}^{\alpha} = \{\beta \in \mathcal{O}, \phi^{S}_{Ax}(\beta) = \phi^{S}_{Ax}(\alpha)\}$, that is, the set of axioms in \mathcal{O} that have the same **structure** as α . We can, finally, define popularity $\pi_{\mathsf{Ax}^{\alpha}}$ of an entity $f \in \mathsf{sig}(\mathcal{O})$ as

$$\pi_{\mathit{Ax}^{lpha}}(f) = rac{|\{eta \in \mathit{Ax}^{lpha}, f \in \mathit{sig}(eta)\}|}{|\mathit{Ax}^{lpha}|}$$

that is, the number of axioms in Ax^{α} that reference f over the size of Ax^{α} itself.

We can plug-in popularity as defined above into a place-holder replacement function and therefore in our distance as follows: When computing a distance between two entities, namely e_1 and e_2 , for each axiom α where either occurs, the function replaces e_1 or e_2 with * and decides whether to replace the other entities with a place-holder depending on their popularity across all the axioms that have the same structure as α .

Definition 5 (Popularity based place-holder replacement). Let \mathcal{O} be an ontology, $e \in sig(\mathcal{O})$ an entity, and $\alpha \in \mathcal{O}$ an axiom. Let Ax^{α} and $\pi_{Ax^{\alpha}}$ be respectively the set of axioms sharing the same structure as α and the popularity metric defined in Definition 4. Finally, let σ be a function that we call **popularity criterion** and maps a popularity value into the set {true, false}.

```
\forall f \in \mathit{sig}(\mathcal{O}), \ we \ define \ our \ function \ as \ follows: \phi_e^{\alpha}(f)
```

$$\begin{array}{ll} - & *if \ f = e; \\ - & f \ if \ \sigma(\pi_{\mathsf{A}\mathsf{X}^\alpha}(f)) = \mathsf{true}; \\ - & \phi^S(f) \ otherwise. \end{array}$$

We can now use the popularity based place-holder replacement defined above in our distance (Definition 3). Given two entities e_1 and e_2 according to the formula we need to compute $\mathsf{Ax}_{\phi}(e_1)$ and $\mathsf{Ax}_{\phi}(e_2)$. For every axiom α in the ontology $\mathcal O$ that references e_1 (resp. e_2), we compute $\phi_{Ax}(\alpha) = \phi_{e_1}^{\alpha} A_x(\alpha)$ (resp. $\phi_{Ax}(\alpha) = \phi_{e_2}^{\alpha} A_x(\alpha)$). Informally, for each axiom, we compute our replacement function based on the popularity of the entities across the set of axioms sharing the same structure as the axiom we are currently considering. In the definition above we deliberately parameterised the decision criterion to make our distance framework independent from any particular implementation. In this work, however, we compute a confidence interval [l,u] for the mean value of $\pi_{\mathsf{Ax}^{\alpha}}$. (95% confidence). We assume the variance is unknown; therefore in order to compute the area under the distribution function (z), we use the values for the t distribution, rather than the normal one in the formulas:

$$l = M - z \cdot \frac{\operatorname{sd}}{\sqrt{N}}, \quad u = M + z \cdot \frac{\operatorname{sd}}{\sqrt{N}}$$

where with sd we denote the standard deviation and with M the mean computed on the set of entities (whose size is N) in the ontology. If the popularity of a given entity is greater than u then we assign true to our σ (see Definition 5), false otherwise.

Example 3. Once again, let our ontology be the Pizza ontology and let us use as our place-holder replacement function ϕ , the one in Definition 5 (based on popularity). Let us compute the replacements for the same axioms as in Example 2. We omit the calculations but the confidence interval for the popularity when applied to such axioms are such that the only entities which will not be replaced are: hasTopping and TomatoTopping, therefore:

The extensive usage of object property has Topping in this particular kind of axiom is the reason why our place-holder replacement function deems it as important and preserves it in the replacement result.

We observe, however, that deciding replacements based on confidence intervals is strongly dependant on the quality of the sample data. TomatoTopping, for instance, in the example above, is judged *popular* too. The reason is that all pizzas in the ontology have TomatoTopping (and MozzarellaTopping) among their toppings. Conversely, the formula correctly spots that several other entities (Margherita, Pizza, hasBase, ...) are not *relevant* when dealing with axioms presenting a particular structure (owlClass *SubClassOf* owlObjectProperty **some**

owlClass). We claim that this is preferable w.r.t. making an *a priori* decision, maybe based on users' intuitions, on what should be replaced and when.

Agglomerative hierarchical clustering: To complete the discussion of our implementation for Algorithm 1, we need to illustrate how we update the distances in our proximity matrix at every agglomeration and what we use as our stopping criterion. For the former we use the Lance-Williams formula (see Section 8.3.3 in [13] - page 524). This formula computes the distance between cluster Q and R, where R is the result of a merger between clusters A and B, as a function of the distances between Q, A, and B. The distance between two sets (clusters) is a function of the distance between their single elements. There are several approaches to compute this, each corresponds to a different value configuration of the coefficients in the general Lance-Williams formula. In the experiments described in the following sections, we used the so-called centroid configuration⁵.

As its stopping criterion, our implementation uses a heuristic decision:

Definition 6 (Agglomerate decision function). Let \mathcal{O} be an ontology and d a distance function. We define the function $\mathsf{agg}_d: 2^{\mathsf{sig}(\mathcal{O})} \times 2^{\mathsf{sig}(\mathcal{O})} \to \{\mathsf{true}, \mathsf{false}\}$ as follows: Given $E = \{e_1, e_2, \dots, e_n\}$ and $F = \{f_1, f_2, \dots, f_m\}$ be two clusters, $\mathsf{agg}_d(E, F) =$

- false if $\exists 1 \leq i \leq n (\exists 1 \leq j \leq m : d(e_i, f_i) = 1);$
- true otherwise.

The algorithm terminates when no pair in the current set of clusters returns true for the Agglomeration decision function agg_d , defined above. When clustering the Pizza ontology, our implementation returns 17 clusters containing over 110 entities in total; these include: a cluster for the toppings that are used in pizzas; one for the named pizzas (pizza with a name and a description of their toppings); and one for the country of origin of the toppings.

As intuitive these groups may seem, given the average familiarity people have with the pizza domain, this represents a cluster analysis based on the actual usage of the entities in the ontology. In this example clusters seem to follow the taxonomy quite well, however, as we shall see in the next section this may not be the case. Performing this kind of analysis can indeed reveal common use between entities that are far apart in the taxonomical hierarchy.

Description of the clusters: Once the clusters are available, the axioms that reference entities in the same cluster can be *generalised* and provide a more abstract view on the entire cluster. We can define a generalisation as a simple substitution of an entity with a variable within an axiom. More formally

Definition 7 (Generalisation). Let \mathcal{O} be an ontology, $E = \{e \in sig(\mathcal{O})\}$ a set of entities, and $\alpha \in \mathcal{O}$ an axiom. Let us now choose a symbol (variable name),

⁵ Although vaguely related, not to be confused with a *centroid* in the K-MEANS cluster analysis - see Chapter 8 in [13].

say ?x. We generalise over E with ?x in α ($g(\alpha, E, ?x)$) when we replace every element of E in α with ?x.

In the definition above, as well as in the remainder of the paper, we will borrow the syntax for variables from OPPL⁶, a declarative language for manipulating OWL ontologies [6]. However, for the purpose of this paper, it is sufficient to say that an OPPL variable can be: **Input or Non generated**, i.e.: they can replace entities in axioms of the corresponding type (there are OPPL variables for each type of entity in a signature); **Generated**, i.e. their value is the result of an expression depending on other variables.

Example 4 (Generalised Pizzas). Let \mathcal{O} be our Pizza ontology and let cluster₁ be the cluster of all the toppings used in pizzas obtained using our cluster analysis above, and cluster₂ be the cluster of all pizzas. Given $\alpha = \mathsf{Margherita} \ \mathsf{SubClassOf}$ has Topping some Tomato Topping:

- $-g(\alpha, cluster_1, ?cluster_1) = Margherita SubClassOf hasTopping some ?cluster_1;$
- $-g(\alpha, {\sf cluster}_2, ?{\sf cluster}_2) = ?{\sf cluster}_2 \ {\sf SubClassOf} \ {\sf hasTopping} \ {\sf some} \ {\sf TomatoTopping};$ or composing the two
- $-g(g(\alpha, \mathsf{cluster}_2, ?\mathsf{cluster}_1), \mathsf{cluster}_1, ?\mathsf{cluster}_1) = ?\mathsf{cluster}_2 \ \textit{SubClassOf} \ \mathsf{hasTopping} \ \mathbf{some} \ ?\mathsf{cluster}_1$

where ?cluster₁ and ?cluster₂ are two variables of type class. (In OPPL: ?cluster₁ :CLASS, ?cluster₂:CLASS).

Generalisations provide a synthetic view of all the axioms that contribute to generate a cluster of entities. Each of these axioms can indeed be regarded as an instantiation of a generalisation, as they can be obtained by replacing each variable in g with entities in the signature of the ontology.

3 Results and Evaluation

Four ontologies (AminoAcid 7 , OBI 8 , a module of the SNOMED-CT 9 containing axioms about hypertension and KUPO 10) were selected for testing the clustering framework; Table 1 summarises some results.

The AminoAcid ontology has been developed internally, allowing comments on the clusters from the ontology's authors. The remaining ontologies are documented, enabling further analysis and evaluation of the clusters and regularities. All of the selected ontologies preexisted the clustering framework. A quantitive analysis was also performed on 85 ontologies from the BioPortal¹¹ repository for

⁶ http://oppl2.sourceforge.net

⁷ http://www2.cs.man.ac.uk/ mikroyae/2011/

iswc/files/amino-acid-original.owl

⁸ http://purl.obolibrary.org/obo/obi.owl

http://www2.cs.man.ac.uk/ mikroyae/2011/iswc/files/ sct-20100731-stated_Hypertension-subs_module.owl

¹⁰ http://www.e-lico.eu/public/kupkb

http://bioportal.bioontology.org/

Ontology name	No of Clusters	No of Clustered entities	Cluster coverage per
			generalisation $(\%)$
AminoAcid	16	77 (84%)	78
OBI	445	2832 (70%)	70
KUPKB	26	470 (44%)	51
SNOMED-CT	77	420 (80%)	62

Table 1. Clustering results on the four selected ontologies

which clusters were computed in less than 3 minutes. The framework was tested with the asserted information of the ontologies, for analysing the regularities of their construction.

The number of clusters in Table 1 shows that in most cases more than 50% of the entities in the ontologies were clustered. In principal, the union of the generalisations describes the cluster, thus a single generalisation might not be necessarily applicable for all the values in the cluster. However, in all four ontologies the majority of the values in a cluster is covered by a single generalisation (cluster coverage percentiles in Table 1).

The results in all of the ontologies are of similar impact and can be found in detail online¹². In this section, however, we will highlight some cases from each ontology.

Inspecting regularities: We detected regularities in all four ontologies and we evaluate them by referring to their documentation or ontology authors. In OBI ontology, ?cluster₇ includes 47 T cell epitopes of specific type that are equivalent classes. The structure of the definition of these terms is similar. To demonstrate this regularity we will consider two classes from ?cluster₇, the "epitope specific killing by T cells" and the "epitope specific T cell activation". These classes appear on the left hand side of the axioms:

```
\alpha = 'epitope specific killing by T cells' EquivalentTo 'T cell mediated cytotoxicity' and ('process is result of' some 'MHC:epitope complex binding to TCR') \alpha = 'epitope specific T cell activation' EquivalentTo 'T cell activation' and ('process is result of' some 'MHC:epitope complex binding to TCR')
```

The generalisation for these axioms is:

```
g(\alpha) = {\sf cluster}_7 \; {\sf EquivalentTo} \; {\sf ?cluster}_8 \; {\sf and} \; ({\sf 'process is result of' some cluster}_{91}) \ where \; {\sf ?cluster}_8 = \{{\sf 'T cell mediated cytotoxicity', 'T cell activation}\}, \ {\sf ?cluster}_{91} = \{{\sf 'MHC:epitope complex binding to TCR'}\}
```

?cluster₈, ?cluster₉₁ are placeholders for the corresponding classes, while the object property 'process is result of' does not belong to any cluster, thus it is not represented by a placeholder.

In [10], a methodology for developing 'analyte assay' terms in OBI using a template based on spreadsheets is described. The clustering framework, gave such a

http://www2.cs.man.ac.uk/~mikroyae/2011/iswc/

cluster of classes (cluster₃₅) and their descriptions in the form of generalisations highlighting their commonalities. For example, there are 13 axioms in cluster₃₅ covered by the following generalisation, which describes the analyte assays that are used to "achieve a planned objective":

```
g(\alpha) = ?cluster<sub>35</sub> SubClassOf ?cluster<sub>117</sub> some ?cluster<sub>16</sub>
```

 $\textbf{Example instantiation: } \alpha = \text{'genotyping assay'} \ \textit{SubClassOf} \ \text{achieves_planned_objective'}$ some 'sequence feature identification objective'

In [8] the design process of the KUP ontology is explained and two main patterns are described for generating the cell types in the ontology. The results of the clustering framework showed such clusters of cells and clusters of classes used as fillers of the properties describing the cells (e.g. participates_in, part_of). Two example generalisations capturing these regularities are:

```
\begin{aligned} &1.\ g(\alpha) = ?\mathsf{cluster}_{13}\ \textit{EquivalentTo}\ ?\mathsf{cluster}_{27}\ \mathsf{and}\ (\mathsf{part\_of}\ \mathsf{some}\ ?\mathsf{cluster}_2), \\ &where\ ?\mathsf{cluster}_{13}, ?\mathsf{cluster}_{27}, ?\mathsf{cluster}_2: \mathsf{CLASS},\ ?\mathsf{cluster}_{27} = \{\mathsf{cell}\} \end{aligned}
```

Example Instantiation:

- $\alpha =$ 'bladder cell' *EquivalentTo* cell **and** (part_of **some** 'bladder')
- 2. $g(\alpha) = ?$ cluster₁ SubClassOf (participates_in some ?cluster₁₆) and (participates_in some ?cluster₁₉), where ?cluster₁₆, ?cluster₁₉ : CLASS

Example Instantiation:

```
\alpha='kidney interstitial fibroblast' SubClassOf (participates_in some 'cytokine production') and (participates_in some 'extracellular matrix constituent secretion')
```

Each one of these generalisations corresponds to a different cluster in the ontology. Also, these regularities were described in [8]. The first regularity is encapsulated in the description of the first pattern and the second regularity is encapsulated in the description of the second pattern. Additional regularities were also detected that refer to longer conjunctions of the previous generalisations (e.g. a conjuction of part_of relationships on the right hand side of the axiom).

In addition, the clustering framework gave an alternative view based on the similar usage of the entities in the ontology. It could selectively reveal repeating structures in the ontology that were more difficult to inspect manually. For example, in SNOMED the highlighted classes of Figure 1(a) are grouped in the same cluster and their similar definition is given by the generalisation of Figure 1(b). However, the inspection of regularities through navigation in the class hierarchy is not an easy task because of the high level of nesting and complexity of the hierarchy. The form of regularity of Figure 1(b) is also described in the technical guide of the ontology [1](section 17.2.2., page 180).

Fully covered generalisations: As it has been shown in Table 1 a single generalisation does not necessarily apply to all of the values in the cluster.

However, there are cases that an individual generalisation can be applicable to all values of its cluster. Such an example is shown in Figure 2, taken from

Generalisation:

```
©'Procedure by site (procedure)'
     'Intracranial procedure (procedure)'

    9'Procedure on abdomen (procedure)

        'Abdominal cavity operation (procedure)
    'Procedure on body system (procedure)'

    9'Procedure on cardiovascular system (procedure)

       'Cardiovascular operative procedure (procedure)'
     © 'Procedure on artery (procedure)'
      'Procedure on blood vessel (procedure)
  © 'Procedure on endocrine system (procedure)'
       'Operation on endocrine system (procedure)
  ©'Procedure on genitourinary system (procedure)'
  ©'Procedure on head AND/OR neck (procedure)'
  © 'Procedure on nervous system (procedure)'
    Procedure on organ (procedure)
  ©'Procedure on soft tissue (procedure)'
  © 'Procedure on thorax (procedure)'
       'Surgical procedure on thorax (procedure)'
        'Thorax implantation (procedure)'
  ▼ @ 'Procedure on trunk (procedure)'
        Operation on trunk (procedure)

    ©'Procedure on mediastinum (procedure)'

         Operation on mediastinum (procedure)

    © 'Procedure on heart (procedure)'

    Operation on heart (procedure)

                'Ventricular operation (procedure)'

 © 'Procedure on pelvis (procedure)'
```

(a) Classes of ?cluster₁₈ as shown in Protégé class hierarchy view. The similar entities do not appear to be siblings in the class hierarchy.

```
g(\alpha) = ?cluster<sub>18</sub> EquivalentTo ?cluster<sub>4</sub>
        and ?cluster6
         and (RoleGroup some
          ((?cluster<sub>10</sub> some ?cluster<sub>2</sub>)
          and (?cluster<sub>74</sub> some ?cluster<sub>16</sub>)))
where
?cluster<sub>6</sub>, ?cluster<sub>18</sub>, ?cluster<sub>4</sub>, ?cluster<sub>2</sub>,
?cluster<sub>16</sub> : CLASS
?cluster<sub>10</sub>, ?cluster<sub>74</sub>: OBJECTPROPERTY.
Example Instantiation:
\alpha = \mathsf{SCT\_79537002} EquivalentTo \mathsf{SCT\_118681009}
        and SCT_387713003
         and (RoleGroup some
          ((SCT_260686004 some SCT_129284003))
          and (SCT_363704007 some SCT_113331007)))
SCT_79537002: 'Operation on endocrine
        system (procedure)'
SCT_118681009: 'Procedure on
        endocrine system (procedure)'
SCT_387713003: 'Surgical procedure (procedure)'
SCT_260686004: 'Method (attribute)'
```

(b) Generalisation and example instantiation of ?cluster₁₈

SCT_129284003: 'Surgical action (qualifier value)'

SCT_113331007: 'Procedure site (attribute)'

SCT_113331007: 'Structure of endocrine system (body structure)'

Fig. 1. View and description of entities of cluster₁₈ in SNOMED hierarchy

the AminoAcid ontology. The first generalisation covers more than one axiom corresponding to a single entity in ?cluster1 (an example instantiation is shown in Figure 2). All the amino acids in the ontology are described following the same template, expressed by the set of 4 generalisations of Figure 2, abstracting 202 axioms. This is a clear indication of the impact the abstraction can achieve when trying to comprehend the ontology. The information that describes these entities is gathered in one place and expressed in a synthetic and meaningful way. That is, because each variable represents a cluster of entities. For example, the ontology engineer by looking the instantiation of the first generalisation understands that cluster5 holds all the physicochemical properties of the amino acids and cluster2 holds all the fillers of these properties.

An analysis of the generalisation coverage and axiom coverage for 85 ontologies from BioPortal is presented in Figure 3(a) and Figure 3(b) respectively. The results of Figure 3 show that there are ontologies, which have cluster coverage higher than 30% and in many cases the average number of instantiations per generalisation exceeds 20. It should be marked that most of the ontologies that have

Values:

?cluster₁ : CLASS

?cluster1 = {Alanine, Arginine, Aspartate, Cysteine, Glutamate, Glutamine, Histidine, Isoleucine, Leucine, Lysine, Methionine, Phenylalanine, Proline, Serine, Threonine,

Tiny Aromatic Amino Acid, Tryptophan,

Tyrosine, Valine, Glycine}

Generalisations:

1.?cluster₁ SubClassOf ?cluster₅ some ?cluster₃

2.?cluster $_1.$ IRI?cluster $_7$ " constant"

3.?cluster₁ SubClassOf AminoAcid

 $4. \textit{DisjointClasses:'} set(?cluster_1.VALUES)'$

 $where \ ? cluster_3 : CLASS,$

?cluster₇: ANNOTATIONPROPERTY, ?cluster₅: OBJECTPROPERTY

Example Instantiations:

for the value ?cluster $_1 = \{Alanine\}$ 1.Alanine SubClassOf hasSize some Tiny
Alanine SubClassOf hasSideChainStructure

some Aliphatic

Alanine *SubClassOf* hasCharge **some** Neutral Alanine *SubClassOf* hasPolarity **some** Non-Polar Alanine *SubClassOf* hasHydrophobicity

some Hydrophobic

2. Alanine label "Alanine"

3. Alanine SubClassOf AminoAcid

4. Disjoint Classes: Alanine, Cysteine, Aspartate,

Glutamate, Phenylalanine, Glycine, Histidine,

Isoleucine, Lysine, Leucine, Methionine,

Asparagine, Proline, Glutamine, Arginine, Serine, Threonine, Valine, Tryptophan, Tyrosine

Fig. 2. Values, generalisations and example instantiation of ?cluster₁ in the AminoAcid Ontology

a high cluster coverage percentile they also have a high average number of instantiations per generalisation. For example, in ontology 34 the cluster coverage per generalisation is 87% and the average number of axioms per generalisation is 560. These cases show that very few generalisations can summarise big number of axioms and can give an inclusive description of the clusters in the ontology.


Inspecting irregularities: In the example of Figure 2, we notice that a possible value of ?cluster₁ is the TinyAromaticAminoAcid. This value is covered only by the third generalisation. The axioms describing this class are:

- 1. TinyAromaticAminoAcid EquivalentTo AminoAcid and hasSize some Tiny
- 2. TinyAromaticAminoAcid SubClassOf AminoAcid


The second axiom is redundant causing the TinyAromaticAminoAcid class to result in the same cluster with the amino acids. By removing this axiom, the TinyAromaticAminoAcid no longer is a value of ?cluster₁. This irregularity is a design defect.

However, there were cases that entities were not included in a cluster because their description was a deliberate exception in the regularity. E.g., in SNOMED, the 'Surgical insertion - action (qualifier value)' is not included in cluster₁₆ as it is not used in a regular axiom like the 'Transplantation - action (qualifier value)':

 $\alpha=$ 'Transplantation (procedure)' *EquivalentTo* 'Procedure by method (procedure)' **and** (RoleGroup **some** ('Method (attribute)' **some** 'Transplantation-action (qualifier value)')), $g(\alpha)=$?cluster₃₀ *EquivalentTo* 'Procedure by method (procedure) **and** (RoleGroup **some** (?cluster₇₄ **some** ?cluster₁₆)))



(a) average cluster coverage per generalisation



(b) average instantiations per generalisation

Fig. 3. Graph showing selected clustering results for 85 ontologies in BioPortal

Values:

 $? \mathsf{cluster}_6 : \mathsf{CLASS}$

?cluster₆ = {SmallHydrophilicAminoAcid,

Small Hydrophobic Amino Acid,

 ${\sf Small Non Polar Amino Acid}, {\sf Small Polar Amino Acid},$

 ${\sf SmallPositiveAminoAcid}\},$

Generalisation:

 $g(\alpha) = ?$ cluster $_{9}$ Equivalent Equiv

where clustery, clustery. CEASS

Example instantiation:

α = SmallPositiveAminoAcid EquivalentTo
 PositiveChargedAminoAcid and SmallAminoAcid

Fig. 4. Values, generalisations and example instantiation of ?cluster₆ in the AminoAcid Ontology

Values:

?cluster₈ = {TinyHydrophobicAminoAcid,

 $TinyNonPolarAminoAcid,\ TinyPolarAminoAcid\}$

Generalisation:

 $g(\alpha) = ?$ cluster₈ EquivalentTo AminoAcid and (?cluster₅ some ?cluster₁₀)

and (?cluster₅ some ?cluster₃)

 $where ?cluster_3 = {Tiny (CLASS)}$

?cluster₅ : OBJECTPROPERTY,

?cluster $_{10}$: CLASS.

Example instantiation:

 $lpha_2={\sf TinyPolarAminoAcid}\ {\it EquivalentTo}$ AminoAcid and hasSize some ${\sf Tiny}$ and hasPolarity some ${\sf Polar}$

Fig. 5. Values, generalisations and example instantiation of ?cluster₈ in the AminoAcid Ontology

To evaluate the sensitivity of the framework, we edited regular axioms of the ontologies to check if these can be identified by the algorithm. In the AminoAcid ontology these irregularities already existed. In particular, cluster₆ and cluster₈ include equivalent classes in the ontology, which are used to categorise the amino acids according to their chemical properties. The equivalent classes are grouped into different clusters according to their regular usage and description. In Figures

4, 5 the description of cluster₆ and cluster₈ is presented respectively. Cluster₆ includes small amino acids with an additional chemical property. On the other hand, cluster₈ has a different design from this of cluster₆ (e.g. it would be expected to be α =TinyPolarAminoAcid EquivalentTo TinyAminoAcid and PolarAminoAcid). For preserving a regular design in the ontology, the terms of cluster₈ are transformed similar to the design of terms of cluster₆. This change has an impact on the clustering; the number of clusters is decreased, but these are more homogenous.

For the rest of the three ontologies, we removed 2 axioms that appear to have a regularity and check if the clustering framework corresponded to these changes. All the entities of which regular axioms were modified were discarded from their initial cluster because the regularity no longer existed.

4 Conclusions and Future Work

We presented a framework for identifying regularities and clustering the entities in an ontology according to these regularities. The application of the approach to 4 ontologies detected regularities that were expected to be found. These findings we confirm and evaluate through access to the ontology authors or the published documentation from the ontology's authors. The framework also provided an alternative presentation of an ontology based on its regularities, giving an insight about the major components of its construction. The generalisations provided a meaningful abstract form of axioms, in which each variable was representing a cluster of similar entities. This abstraction has potential as a tool for comprehension of an ontology, as it manages to summarise axioms in a coherent way. The method tended to give good coverage of an ontology's axioms within clusters, suggesting that it could show the significant portions of the ontology. The analysis on the BioPortal ontologies also highlighted such cases. That not all axioms are clustered is also meaningful; not all axioms can be part of regularities and those that do not cluster can indicate deviations from a style or simply deliberate authoring techniques. Either reason could be informative to a person comprehending an ontology.

The evaluation showed that changing regularities affected clustering results. Altering referencing axioms of entities that belonged to a cluster either caused them to be regrouped in a different cluster or their exclusion from any cluster. This shows the method to be sensitive to changes and a potential tool for helping authors to "tidy up" their ontology. Future work will include the development of plugins for the Protégé editor for supporting the development process dynamically through access to these features.

The inspection of regularities in known ontologies helped us to derive knowledge patterns when they existed. We expect that the further exploitation of the generalisation forms and the examination of possible combinations of these can lead to the induction of patterns in the ontology (e.g. like the pattern describing the amino acids in Figure 2). Future work will also involve alternative decisions of a transformation policy and of a clustering algorithm. Although the current implementation, which is based on a popularity transformation, produced adequate results, it will be worth examining other techniques that can generate

homogenous and well defined clusters. The inspection of regularities in known ontologies helped us to derive knowledge patterns when they existed.

Our use of a basic clustering approach has a demonstrable use in finding regularities and irregularities in an ontology. It has potential for offering authors a means to gain generalisation of the major portions of an ontology; to detect deviations from a given style of representation and to facilitate the comprehension of what can be large and complex logical artefacts. As ontologies using OWL can be large and complex, the provision of techniques to manage this complexity, especially when attempting to understand an ontology's construction, should be an important addition to an ontology author's toolbox.

References

- Snomed technical reference guide (January 2011), http://www.ihtsdo.org/fileadmin/user_upload/Docs_01/ Publications/SNOMED_CT
- Egaña, M., Rector, A., Stevens, R., Antezana, E.: Applying ontology design patterns in bio-ontologies (2008)
- 3. Gangemi, A.: Ontology Design Patterns for Semantic Web Content. In: Gil, Y., Motta, E., Benjamins, V.R., Musen, M.A. (eds.) ISWC 2005. LNCS, vol. 3729, pp. 262–276. Springer, Heidelberg (2005)
- 4. Gangemi, A., Presutti, V.: Ontology design patterns. In: Handbook on Ontologies, pp. 221–243. Springer, Heidelberg (2009)
- Horridge, M., Parsia, B., Sattler, U.: Laconic and Precise Justifications in OWL. In: Sheth, A.P., Staab, S., Dean, M., Paolucci, M., Maynard, D., Finin, T., Thirunarayan, K. (eds.) ISWC 2008. LNCS, vol. 5318, pp. 323–338. Springer, Heidelberg (2008)
- Iannone, L., Aranguren, M.E., Rector, A.L., Stevens, R.: Augmenting the expressivity of the ontology pre-processor language. In: OWLED. CEUR Workshop Proceedings, vol. 432. CEUR-WS.org (2008)
- Iannone, L., Rector, A., Stevens, R.: Embedding Knowledge Patterns into OWL. In: Aroyo, L., Traverso, P., Ciravegna, F., Cimiano, P., Heath, T., Hyvönen, E., Mizoguchi, R., Oren, E., Sabou, M., Simperl, E. (eds.) ESWC 2009. LNCS, vol. 5554, pp. 218–232. Springer, Heidelberg (2009)
- 8. Jupp, S., Horridge, M., Iannone, L., Klein, J., Owen, S., Schanstra, J., Stevens, R., Wolstencroft, K.: Populous: A tool for populating ontology templates. Arxiv preprint arXiv:1012.1745 (2010)
- 9. Parsia, B., Sirin, E., Kalyanpur, A.: Debugging OWL ontologies. In: Proceedings of the 14th International Conference on World Wide Web, pp. 633–640. ACM (2005)
- Peters, B., Ruttenberg, A., Greenbaum, J., Courtot, M., Brinkman, R., Whetzel, P., Schober, D., Sansone, S., Scheuerman, R., Rocca-Serra, P.: Overcoming the ontology enrichment bottleneck with quick term templates (2009)
- Šváb-Zamazal, O., Svátek, V.: Analysing Ontological Structures Through Name Pattern Tracking. In: Gangemi, A., Euzenat, J. (eds.) EKAW 2008. LNCS (LNAI), vol. 5268, pp. 213–228. Springer, Heidelberg (2008)
- 12. Šváb-Zamazal, O., Svátek, V., Iannone, L.: Pattern-Based Ontology Transformation Service Exploiting OPPL and OWL-API. In: Cimiano, P., Pinto, H.S. (eds.) EKAW 2010. LNCS, vol. 6317, pp. 105–119. Springer, Heidelberg (2010)
- 13. Tan, P.-N., Steinbach, M., Kumar, V.: Introduction to Data Mining. Addison-Wesley (2005)