

Reasoning on Crowd-Sourced Semantic Annotations to Facilitate Cataloguing of 3D Artefacts in the Cultural Heritage Domain

Chih-Hao Yu, Tudor Groza, and Jane Hunter

School of ITEE, The University of Queensland, Australia
chih.yu@uqconnect.edu.au, {tudor.groza,j.hunter}@uq.edu.au

Abstract. The 3D Semantic Annotation (3DSA) system expedites the classification of 3D digital surrogates from the cultural heritage domain, by leveraging crowd-sourced semantic annotations. More specifically, the 3DSA system generates high-level classifications of 3D objects by applying rule-based reasoning across community-generated annotations and low-level shape and size attributes. This paper describes a particular use of the 3DSA system – cataloguing Greek pottery. It also describes our novel approach to rule-based reasoning that is modelled on concepts inspired from Markov logic networks. Our evaluation of this approach demonstrates its efficiency, accuracy and versatility, compared to classical rule-based reasoning.

1 Introduction

Advances in 3D data acquisition, processing and visualisation are providing cultural institutions with novel methods for preserving cultural heritage and for making it more accessible to scholars and the public, via online search interfaces. An increasing number of museums are using 3D scanning techniques to overcome the limitations of 2D data representations and to improve access to high quality surrogates of fragile and valuable artefacts via the Internet [1–3].

The growing popularity of scanning of cultural heritage artefacts to generate 3D surrogates, has led to the need for online repositories, to store these large scale collections and for tools and user interfaces to make these resources easily accessible, sharable and retrievable [4]. However, as the size of the online collections expands, the ability to catalogue and search such collections becomes increasingly difficult [5]. Museums are finding the cost of providing rich contextual metadata for such collections prohibitive and hence are keen to explore other ways to generate it, for example, by exploiting social tagging and annotation services [6].

Several successful examples of applying Web 2.0 techniques in the museum context have been provided in the recent past, most of which focus on social tagging for 2D images, e.g., the Powerhouse Museum OPAC2.0 project¹ [7] and the

¹ <http://www.powerhousemuseum.com/museumexchange/>

Steve.Museum project [5]. A list of successful case studies of crowdsourcing in the cultural heritage domain is presented in the work of Oomen and Aroyo [8]. In this paper we present an innovative application – the 3DSA system – which takes social tagging and semantic annotations to a new level, by applying them to the rich and interactive environment of 3D digital surrogates of cultural heritage artefacts. In particular, we focus on leveraging crowd-sourced semantic annotations to streamline the cataloguing of 3D museum artefacts via probabilistic reasoning, which will lead to improved search and retrieval of 3D objects.

Traditionally, museum artefacts are manually curated and classified by experts using a process of literature search and comparison with similar related artefacts. This is a laborious and time-consuming task. However certain classes of objects in the cultural heritage domain are governed by well-defined rules and guidelines (e.g., Greek vases). Furthermore, some of the attributes used in these guidelines can be automatically acquired (e.g., spatial dimensions and shapes of artefacts), while others require human interpretation (e.g., the characters depicted on the surface of the artefacts). Hence, our hypothesis is that: given a set of domain-specific rules (documented by the domain experts) for classifying an artifact, together with a set of artefact properties/attributes that have been either manually or automatically acquired, the (high-level semantic) classification task can be automated by performing rule-based reasoning. More specifically, assuming curatorial/anthropological experts define the rules (for a particular class of artefacts), and cultural heritage scholars, students and enthusiasts semantically annotate components of the artefacts, then the system can combine these manual tags with automatically extracted size and shape attributes, to provide on-demand the most likely classification for given 3D objects.

This is a key goal and the underlying methodology of our 3DSA system. Using a custom-built, domain-specific ontology, the 3DSA system enables semantic annotation of points, surface regions and 3D volumetric segments of 3D objects. Users can interactively specify the low-level features and annotate them with tags drawn from ontology. These annotations are then used by a rule-based reasoning engine to infer high-level descriptions of the 3D objects, which in turn, facilitate the classification and cataloguing of the 3D artefacts. In this paper, we describe the solution we have developed to perform this process by adopting a Markov logic network approach and combining first order logic rules with probabilities. The evaluation results indicate that our approach outperforms classical first order logic in terms of both efficiency and accuracy and at the same time attracts excellent feedback in usability tests.

The 3DSA system is deployed in the R D Milns Antiquities museum at The University of Queensland and is used by a wide range of users, from postgraduate students and museum staff to educators and archeologists. The current application focus is Greek pottery, both because this topic is a major interest of the museum and also because our current 3D scanner can easily accommodate objects of this size.

The remainder of the paper is organised as follows. Section 2 presents the high level overview of the system architecture, while Section 3 presents a brief

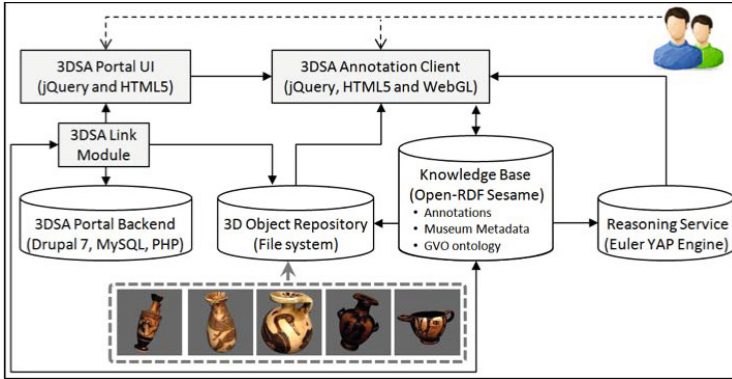


Fig. 1. Overview of the 3DSA system architecture

description of our use case and its associated inference and retrieval requirements. In Section 4, we describe in detail the probabilistic reasoning mechanism. Before concluding in Section 7, we discuss some experimental results (Section 5) and the existing related work (Section 6).

2 3DSA System Overview

The 3DSA system provides a Web interface that enables users to explore collections of 3D objects via a browser interface, as well as select individual 3D objects and attach semantic annotations to either the entire object or to points, surface regions or segments of those 3D objects. Fig. 1 depicts the high-level architecture of the system, which comprises seven major components.

The front-end of the system consists of two clients:

- the 3DSA Annotation client – an Open Annotation (OA) client, developed using jQuery, HTML5 and WebGL with O3DWebGL library, which provides the 3D interaction and semantic annotation capabilities, and
- the 3DSA Web Portal – an online gallery, developed using Drupal 7, jQuery and HTML5, which enables exploration and searching capabilities.

The back-end of the system is composed of:

- the 3D object repository – responsible for storing the 3D objects in their native formats;
- the 3DSA Web portal backend – responsible for serving and managing Web pages and contents for Web portal using Drupal 7 CMS, PHP and MySQL as the underlying database.
- a knowledge base that stores museum metadata, domain specific ontologies and semantic annotations – developed on top of an OpenRDF Sesame server;
- the 3DSA link module – a custom built module for Drupal 7 to provide support for managing and uploading 3D objects, as well as access to the annotation client and the semantic search features offered by the knowledge base, and

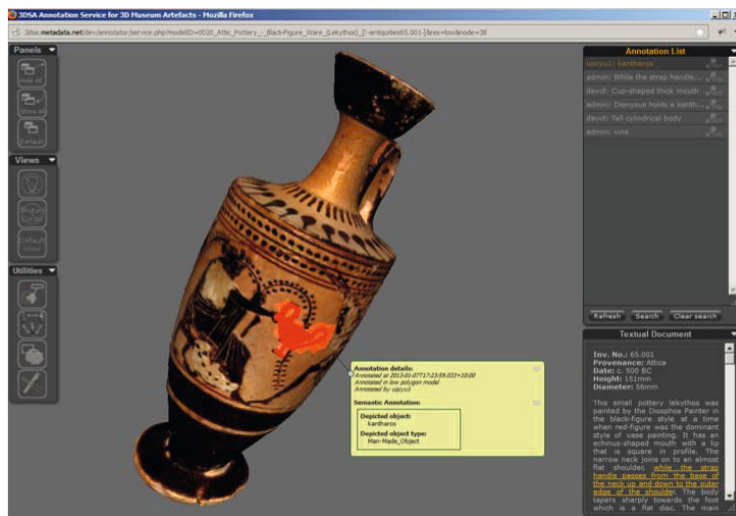


Fig. 2. The main front-end of the 3DSA annotation client

- the Reasoning service (built using the Euler YAP Engine (EYE) reasoner²)
- that provides inference capabilities based on expert-curated rules, crowd-sourced semantic annotations and metadata automatically generated from size and shape measurement.

Fig. 2 illustrates the main user interface of the 3DSA annotation client displaying a 3D representation of an “Attica Black-Figure lekythos” from 500BC, with an annotation attached to a surface depiction of a “kantharos”. 3DSA offers complex manipulation and viewing functionalities, including dynamic viewpoint perspective (i.e., only annotations relevant to the current viewpoint are displayed), viewing objects in different resolutions (high, medium, low) – subject to the available bandwidth and graphical computing power, as well as real-time collaborative annotation visualisation.

To support the current use case of the system, i.e., Greek pottery, we have developed the Greek Vases ontology (GVO) as an extension CIDOC-CRM³ – the *de facto* standard for defining conceptual entities and relationships in the cultural heritage domain. GVO defines the atomic components (parts) of Greek vases (e.g., body, mouth, neck), in addition to the illustrative decorations typically depicted on them (e.g., patterns, inscriptions, figures) and their associated attributes (e.g., sizes, shapes, colour). Excluding the CIDOC-CRM components, the current version of GVO ontology contains 123 concepts, 76 properties and 1534 instances – the collection of instances will expand in time with the increase of the community contribution.

² <http://eulersharp.sourceforge.net/>

³ <http://www.cidoc-crm.org/>

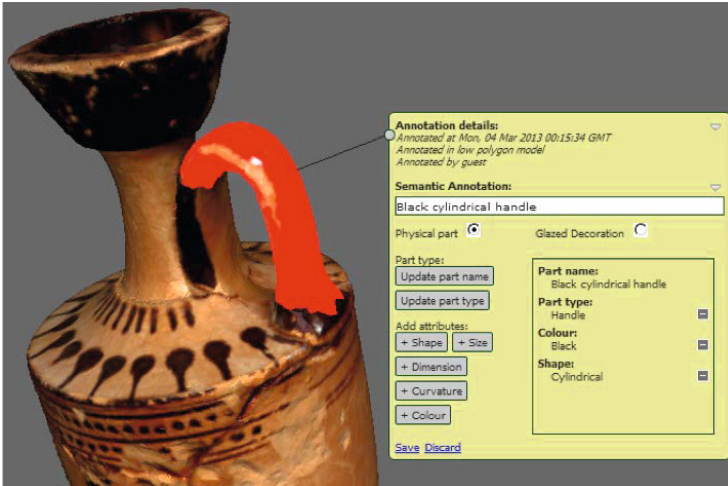


Fig. 3. Semantically annotated region of a Greek vase

In terms of functionality, end-users can interactively specify points and polygon shapes overlaid on the displayed 3D object (to define surface regions and segments) and annotate them with concepts from GVO. Moreover, the system guides the user in the semantic annotation process, by suggesting ontological concepts via auto-completion or via pre-defined structural elements (e.g., part types or attributes). For example, Fig. 3 depicts a fully specified semantic annotation attached to the handle of a 3D Greek vase, which includes the part type (**Handle**), its colour (**Black**), and its shape (**Cylindrical**). The annotation result is serialised using the Open Annotation Data Model⁴ that has been extended via a 3D fragment identifier based on X3D specification [9].

3 Use Case

As described in the Introduction, the 3DSA system is deployed and used in the R D Milns Antiquities Museum at the University of Queensland. The museum's aim is to crowd-source semantic annotations from online scholars, students, cultural heritage enthusiasts and other interested end-users to assist with the cataloguing of their Greek pottery collection. Following a series of discussions with the museum experts, we were able to define usage scenarios and the necessary requirements for the 3DSA system. An example of a practical scenario is presented below.

Collaborative cataloguing scenario. A university educator creates an interactive assessment for his students. The students are required to identify the components of a scanned Greek container and to provide specific descriptions of its parts. Each student is able to create his/her own annotations. For example, the

⁴ <http://www.openannotation.org/spec/core/>



Fig. 4. Ranking of probable shapes for the given 3D object, based on the 3DSA reasoning capability

first student specifies the following attributes of the **body**, e.g., *small* and *deep*, a second student describes the **handles**, e.g., *vertical handle* for one and *horizontal handle* for the other, while a third one defines the **mouth** as being *wide*. The museum curator, in charge of cataloguing the newly scanned object, typically, uses his/her prior knowledge and the literature on Greek pottery shapes to classify this artefact. However, since the students have already collectively annotated the parts of the container, the museum expert is able to leverage this crowd-sourced knowledge to infer a set of candidates. Using the 3DSA inference engine, the curator can retrieve a ranked list of candidates (based on the size and shape of its parts), together with their probability (see Fig. 4). The curator observes that the top result (83.33% probability) is a “*Skyphos Type B*” shape, due to its “*small deep bowl-shaped body*”, “*vertical handle*”, “*horizontal handle*”, “*wide mouth*” and “*thin foot*”. After consulting the literature on “*Skyphos Type B*” vases, the curator confirms the validity of the result and the artefact is tagged/catalogued appropriately. Additionally, based on the underlying ontology, the newly acquired object is also automatically be tagged with the synonyms, “*Glaux Skyphos*”, “*Skyphos*” and “*Drinking cup*”.

Using such scenarios we have gathered the following set of requirements:

Crowd-Sourcing Semantic Annotations. As mentioned, the museum is keen to exploit the power of collective intelligence, and hence the system should enable collaborative semantic annotation of points, surface regions and segments of 3D artefacts.

Flexible Classification of Artefacts. The most laborious task in this context is determining the most reliable classification for a given artefact. The system needs to take advantage of the crowd-sourced semantic annotations and combine them with metadata generated from the 3D scanning process to enable precise cataloguing. Furthermore, the system needs to be flexible and take into account possibly incorrect semantic annotations, hence leading to the need for a soft classification (based on probability), as opposed to a hard multi-class classification.




Fig. 5. Example of a simple rule defining “*Skyphos Type B*”

Defining Expert Rules. The Greek pottery domain is governed by a fairly well-established set of cataloguing guidelines based on the size and shape of their parts⁵. Our approach requires experts to encode these guidelines into rules, which will then drive the inference. Consequently, the rule definition process needs to be simple and oriented towards non-technical users (e.g., as depicted in Fig. 5).

4 Reasoning Capabilities in 3DSA

The expert rules are used to derive high-level artefact descriptions from low-level annotated features. To enable reasoning across the annotations, we are using the Euler YAP Engine (EYE)⁶. The EYE reasoner is a high performance backward-chaining inference engine enhanced with Euler path detection. The rules used for inferencing are formatted in N3Logic – an extension of RDF that enables the usage of the same language for both logic and data. According to their deep taxonomy benchmark⁷, EYE is the most efficient among other existing popular semantic reasoners such as Jena, Pellet and CWM.

In a Web environment, EYE has two components: the EYE client widget and the EYE HTTP server. The former is a Javascript client extension that provides the communication channel with the EYE server, while the latter is an HTTP interface for the EYE reasoning engine.

Fig. 6 illustrates the data flow of rule-based reasoning capability within 3DSA. On the backend side, the EYE reasoner has direct access to the data stored in the knowledge base. When the user requires inference, the knowledge base sends the latest set of rules to the EYE Reasoner, while the 3DSA Annotation Client invokes the EYE Client Widget. This communicates with the reasoner, retrieves the results and provides them to the Annotation Client, which presents them to the user (see Fig. 4). If the user confirms the validity of one of the results, the corresponding inference is recorded in the knowledge base.

From a functional perspective, the goal is to infer the shape or class of a given 3D object based on the crowd-sourced annotations associated with the low-level shape features. For example, as described in Sect. 3, if an object is annotated with *wide mouth*, *small deep body*, *vertical handle* and *horizontal handle*, then we want to infer with a certain level of precision that the object is a *Skyphos Type B*

⁵ <http://www2.ocn.ne.jp/~greekart/vase/shape.html>

⁶ <http://eulersharp.sourceforge.net>

⁷ <http://eulersharp.sourceforge.net/2003/03swap/dtb-2010.txt>

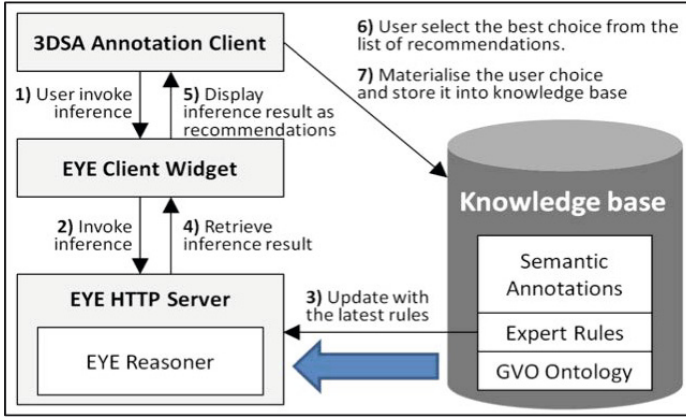


Fig. 6. Data flow in the reasoning component of 3DSA

(*Glaux Skyphos*). Our goal is to present the user with a ranked list of candidates, each of which has an associated probability, instead of a rigid classification result, typical of conventional reasoning. This probabilistic approach is critical in our context, since we rely on crowd-sourced semantic annotations, which may be ambiguous, incomplete or incorrect. In such a context, conventional reasoning will be highly error prone and likely to produce unreliable results.

In order to generate a ranked list of candidates, we combine first order logic with probabilities. More specifically, our solution is based on the concept of *Markov logic network*. Markov logic [10, 11] is a representation formalism that generalises first-order logic (FOL) by attaching probabilities to FOL statements. It defines all unsatisfiable statements as having a probability of zero and all universally true statements as having a probability of one. A collection of such probabilistic statements forms a Markov logic network. Inspired by this framework, we have used the same concept to soften the constraints of FOL – if an entity definition does not satisfy one statement, it is simply returned as less probable rather than impossible.

The difference between the classic Markov logic network and our approach is found in the weight distribution. Instead of using the typical probability distribution as part of a log-linear model, we compute the final probability of a candidate by normalising its sum of weights with the total weight provided by a rule (see Eq. 1).

$$P(A = s) = \frac{1}{\sum_{f_i \in F_s} W_i} \left(\sum_{f_i \in F_s} W_i \prod_{c_j \in C_{f_i}} f_j(g) \right) = \frac{1}{\sum_{f_i \in F_s} W_i} \left(\sum_{f_i \in F_s} W_i G_i \right) \quad (1)$$

The probability P of a 3D artifact A having shape s , is the sum all the weights W_i associated with each feature f_i , within the set of clauses F_s applied on the shape s , multiplied with the grounding G_i of the corresponding feature f_i , normalised by the sum of all weights W_i of shape s . The grounding variable

takes a binary value, $G_i \in \{0, 1\}$, corresponding to a match (1) or not (0) of G_i with the feature f_i .

This approach does not restrict the description of a grounding feature to an atomic element. A grounding feature may very well be expressed via multiple conditions assigned to the antecedent in a rule. The grounding $c_j(g)$ of each corresponding condition c_j within the set C_{f_i} would then belong to a feature f_i that should be considered. The grounding of each condition is expressed, again by a binary variable $c_j(g) \in \{0, 1\}$, with $c_j(g) = 1$ if condition c_j is matched and 0 otherwise. The product of the grounding values for all conditions $\prod_{c_j \in C_{f_i}} f_j(g)$ forms the grounding for the specific feature G_i , which has value of 0 or 1.

Below we present an example of a complex rule, broken down into a set of elementary rules, each of which describes a grounding feature of a Greek pottery shape. In general, the rule has the format listed below:

Shape a gvo:Greek_Pottery_Shape

```
Antecedent{Grounding_feature_1} → Consequent{Shape, ID_1, Weight_1}
Antecedent{Grounding_feature_2} → Consequent{Shape, ID_2, Weight_2}
Antecedent{Grounding_feature_3} → Consequent{Shape, ID_3, Weight_3}
```

Each grounding feature is separately assigned to the antecedent of an individual rule and each rule has a consequent associated with a Greek pottery shape, a unique ID and a weight value. When an antecedent is matched, the shape and its corresponding weight are provided. The reasoner iterates over these results, sums the weights and computes the probability of each Greek pottery shape. Finally, the results are ranked and presented to the user.

To achieve the above using the EYE reasoner, the N3Logic rules for each Greek pottery shape must be written in the following format (the rule below corresponds to the 3D object depicted in Fig. 4):

```
gvo:Skyphos_type_B a gvo:Greek_Pottery_Shape.
?g gvo:hasPart [a gvo:Body; gvo:hasCharacteristic gvo:small].
=> (?g gvo:Skyphos_type_B rule1) ex:giveWeight 0.3.
?g gvo:hasPart [a gvo:Body; gvo:hasCharacteristic gvo:deep].
=> (?g gvo:Skyphos_type_B rule2) ex:giveWeight 0.7.
?g gvo:hasPart [a gvo:Body; gvo:hasCharacteristic gvo:bowl-shape].
=> (?g gvo:Skyphos_type_B rule3) ex:giveWeight 1.0.
?g gvo:hasPart [a gvo:Handle; gvo:attachedAlong gvo:horizontal].
=> (?g gvo:Skyphos_type_B rule4) ex:giveWeight 0.5.
?g gvo:hasPart [a gvo:Handle; gvo:attachedAlong gvo:vertical].
=> (?g gvo:Skyphos_type_B rule5) ex:giveWeight 0.5.
?g gvo:hasPart [a gvo:Mouth; gvo:hasCharacteristic gvo:wide].
=> (?g gvo:Skyphos_type_B rule6) ex:giveWeight 1.0.
?g gvo:hasPart [a gvo:Foot; gvo:hasCharacteristic gvo:thin].
=> (?g gvo:Skyphos_type_B rule7) ex:giveWeight 1.0.
?g gvo:hasMeasurement [gvo:hasHeightValue[math:lessThan 300]].
=> (?g gvo:Skyphos_type_B rule8) ex:giveWeight 1.0.
```

In addition to the specific rules, we have specified the way in which the sum of the weights should be calculated for all Greek pottery shapes/classes, as well as the corresponding probability:

```
{
?TARGET a :Greek_Pottery.
?SHAPE a :Greek_Pottery_Shape.
(?SCOPE ?SPAN) e:findall
  (?WEIGHT
    {(?TARGET ?SHAPE ?RULE) ex:giveWeight ?WEIGHT}
    [math:sum ?PASS_WEIGHT]).
(?SCOPE ?SPAN) e:findall
  (?WEIGHT
    {?P => {(?TARGET ?SHAPE ?RULE) ex:giveWeight ?WEIGHT}}
    [math:sum ?TOTAL_WEIGHT]).
(?PASS_WEIGHT ?TOTAL_WEIGHT) math:quotient ?PROBABILITY.
} => ?TARGET :has_shape (?PROBABILITY ?SHAPE).
```

For example, to determine whether a pot has the shape/class “*Skyphos Type B*” based on the annotations provided by the students in the scenario introduced in Sect. 3, i.e., “*small deep body*”, “*horizontal handle*”, “*vertical handle*” and “*wide mouth*”, the reasoner will sum the weights of the matched statements using the rule given above ($0.3 + 0.7 + 1 + 0.5 + 0.5 + 1 = 5$) and then divide them by the total weight ($0.3 + 0.7 + 1 + 0.5 + 0.5 + 1 + 1 = 6$). This will result in a final probability of 83.33% ($5/6 * 100 = 83.33$) for the given pot to have the shape/class “*Skyphos Type B*”.

5 Evaluation

We have performed several experiments to test both the efficiency as well as the scalability of our reasoning approach. In addition we have conducted a usability study of our cataloguing suggestions approach. Below, we discuss the results of each of these tests. The system used throughout the performance tests was equipped with an Intel Core 2 Duo 2.66GHz CPU, 1.96GB of RAM and an ATI Radeon HD 2400 Pro low-end graphic card. In the case of the usability study, the test users were able to use their own desktop computers.

We ran two experiments to evaluate the reasoning scalability by comparing our probabilistic approach to classical FOL reasoning, using EYE as underlying technology in both experiments. Firstly, we performed reasoning on 50 Greek pottery objects of different shapes and with different sized annotation sets (ranging from 10k to 100k triples). Secondly, we fixed the number of triples to 1 million and generated a diverse range of Greek pottery shape types (1 to 10). Fig. 7 depicts the results of both experiments. As can be observed from the results, our approach has outperformed the FOL reasoning, despite both approaches producing the same reasoning result. In the first case there is a linear dependency between the size of the annotation set for the 3D object and the time required to complete the reasoning. Unfortunately, neither approach achieves tolerable




Fig. 7. Comparison of performance speed: FOL vs. our approach

Table 1. Experimental results achieved by arbitrarily switching statements in the domain rules

Object / rule-set	Original rules		Altered rules	
	FOL	Our approach	FOL	Our approach
1	0.078 s	0s	2.344s	0s
10	0.078s	0.015s	18.422s	0s
20	0.125s	0.016s	27.657s	0.016s
30	0.172s	0.016s	28.813s	0.015s
40	0.172s	0.016s	28.578s	0.015s
50	0.172s	0.016s	28.828s	0.015s

performance, however our approach is faster on average by around 49%. In the second experiment, the difference is much clearer and ranges from double the speed to an improvement of up to two orders of magnitude.

An interesting discovery within this experiment has been that the classical FOL reasoning is sensitive to the ordering of the statements in the rules. We have identified this aspect by arbitrarily switching two statements in the rule that corresponded with each Greek pottery shape (see Table 1). The time required for the FOL inference increases dramatically, even for a small dataset of 500 triples. This is due to EYE’s backward chaining, which is a goal-driven method and therefore sensitive to goal ordering. If a sub-goal in violation does not fail early, it backtracks to all previous sub-goals, which prolongs the time of inference. On the other hand, our approach has not been affected since our complex rules are divided into a set of elementary units, and hence, the amount of backtracking is greatly reduced. This issue will have a significant impact on rule creation if an FOL approach is employed. Museum experts would have to take extra care when ordering the statements that define the rules.

In the second type of experiment, we evaluated the precision of the two approaches by randomly removing annotations attached to different parts of 100 Greek pottery objects and running the inference on the incomplete dataset. The results were remarkable: FOL achieved a precision of merely 19%, while our approach, at $K=1$ (i.e., top result) achieved a precision of 97%, thus proving both

the superiority as well as the versatility of using probabilistic reasoning. The reason behind the poor results of FOL lies in its hard constraint nature, which prevents it from producing any results at all, even when only a single statement in a rule is violated. The correct results are found only because the corresponding matching rules are not affected by the missing triples (removed during the experiment). Our approach, on the other hand, is able to deal successfully with missing information, as each triple removed makes the possible candidates less likely instead of marking them as false negatives. The main drawback of using this versatile inference is that it may produce false positives in the case of closely related rules or when too many annotations are missing (or too many incorrect annotations are provided). Overall, this experiment shows the superiority of using probabilistic reasoning over classical FOL when dealing with semantic annotations with characteristics similar to those obtainable through crowd-sourcing and well-established domain-specific rule sets.

Finally, we performed a usability study with 18 users (postgraduate researchers, museum curators and educators) from a range of disciplines (Anthropology, Archaeology and Classics) at the University of Queensland. The test users were given a set of six 3D objects and asked to manually identify the closest matching shape/vase from a reference set of images. They were then requested to use the 3DSA reasoning capability and select the most suitable shape from the reasoning result. At the completion of those two tasks, the test users were requested to complete a survey and were encouraged to provide verbal feedback.

The survey results indicate that 50% of users found the manual cataloguing of 3D objects to be extremely cumbersome, while all users agreed that the process became very simple once reasoning had been applied. All test users found the reasoning capability of 3DSA to be efficient – 63% very fast and 37% fast. Most of the test users found it helpful to use the 3DSA semantic reasoning to assist the cataloguing of Greek pottery 3D objects – 87% agrees and 13% neutral.

The feedback from the museum staff was highly encouraging. They found 3DSA extremely useful in classifying Greek pottery – especially for museum curators without expertise in Greek pottery. In particular, they appreciated the ranked results, which combined with their personal knowledge and expertise, enabled them to catalogue an item very quickly. On a less positive note, our application was less appealing to anthropologists, as they had doubts about the usefulness of crowd-sourced semantic annotations in classifying ethnographic collections. This is mainly due to the nature of their studies, which do not focus on the individual interpretation of the features and decorations of the artefacts, but rather on the method of creation, provenance and use of such objects.

6 Related Work

The increasing adoption of the Semantic Web into mainstream technologies has led to a wide range of projects that use ontologies and reasoning to annotate and infer new knowledge in the cultural heritage domain (or in other domains that exploit 3D surrogates).

The AIM@SHAPE project aimed at externalising and sharing the knowledge captured within digital shapes via knowledge management techniques applied to shape modelling [12]. The semantics extracted (automatically or semi-automatically) from 3D models has then been used for reasoning and retrieval purposes. One of the projects use cases has been museum collections, e.g., scanned objects from The National Museum of Ethnology. The FOCUS K3D project had similar goals and focused on automatic identification of 3D cultural heritage objects, through reasoning based on manually acquired annotations [13]. Finally, the 3D-COFORM project enables the annotation of 3D cultural heritage objects based on CIDOC-CRM standards [14]. Their Metadata Repository uses OWLIM for rule-based reasoning to provide a more efficient retrieval of 3D objects [14]. The key difference between these approaches and our work, is that we have applied probabilistic reasoning to infer high level semantic tags from low-level features for 3D cultural heritage objects.

From a different perspective, the ArchaeoKM project [15] is a Web based knowledge management system for archaeologists that supports semantic annotations of 3D spatial data and uses ontologies to perform inference on the created and presented knowledge. Objects are identified and tagged with concepts from the domain ontology and rules are then employed to generate new knowledge. Archaeologists can create their own rules to validate or discover knowledge, which is particularly useful for retrieving objects hard to classify but which possess certain features.

Other approaches in the cultural heritage domain, not using 3D surrogates, include: (i) the CULTURESAMPO project – uses Jena to perform reasoning on the explicit descriptions of museum objects with the aim of discovering hidden knowledge; for example, the place of manufacture or the usage of museum objects [16]; (ii) The Gothenburg City Museum – uses BigOWLIM to integrate and reason across multiple datasets (including DBpedia and Geonames, in addition to the PROTON, CIDOC-CRM and MAO ontologies); this results in a rich search and browse experience beyond the specific knowledge externalised by their museum collection [17]; (iii) the CHIP project developed in collaboration with the Rijksmuseum Amsterdam the Art Recommender to infer artwork features (e.g. creator, material, subject) and semantic relations between features (e.g. broader/narrower, style, birth or death place) by taking into account user preferences [6]; and (iv) the MultimediaN E-Culture project, which uses ClioPatria [18] (a framework based on SWI-Prolog and a suite of Semantic Web libraries) to perform reasoning based on rich semantic annotations, in order to facilitate semantic search in large heterogeneous cultural heritage collections [19].

In the other domains, Bucher et al. [20] investigated the application of semantic reasoning to 3D city models in CityGML, to classify buildings in terms of shapes, symmetries and repetition. Similarly, the WiDOP project, designed for the engineering domain [21], exploits reasoning capabilities to semantically annotate 3D objects to match detected geometries with probable objects. Both projects use SWRL rules to perform rule-based semantic inferencing. The ISReal project [22] claims to be the first open and cross disciplinary 3D Internet research

platform for intelligent 3D simulation of realities. It uses OWL-DL, Pellet and the RDF relational reasoner STAR [23] to infer scene descriptions in XML3D using semantic annotations. Finally the Onto3D retrieval system [24] is a research application that investigates retrieval of 3D objects. The project attempts to add semantics to 3D objects by linking low-level features to high-level semantic descriptions from WordNet. It uses SWRL rules to infer semantic properties of 3D objects and stores the results into a general ontology. The ontology is then used to improve both text-based and content-based object retrieval.

Most of the above-described projects use reasoning to enhance content retrieval. The 3DSA system, on the other hand, investigates how cultural heritage artefacts can be enriched through crowd-sourced semantic annotations, in addition to using a completely different approach to conduct reasoning on the resulting data. Furthermore, since we rely on collective knowledge acquisition, we are required to deliver a versatile solution capable of dealing with noisy, ambiguous and incomplete annotation – issues usually disregarded by previous approaches. In terms of 2D compatibility, the dimensional differences of the graphical objects would not influence our reasoning approach, since our rule-based reasoning solution is applicable to both 3D and high-resolution 2D scans. However, the annotation and crowdsourcing components are specifically built to deal with challenges that arise from manipulating 3D objects (e.g., rotate, pan and zoom) and would, implicitly, require a re-design of 3DSA system in order to cater for high-resolution 2D scans. Finally, our approach enables experts to define inference rules in a very simple and straightforward manner, abstracting completely from the underlying reasoning framework.

The 3DSA system is a Web-based application for collaborative semantic annotation of parts and regions of 3D museum artefacts, similarly to native applications such as Tagg3D [25] and 3D-COFORM Integrated Viewer/Browser [26]. However, as far as we are aware, it is also the only application in this domain to apply probabilistic reasoning on 3D segment-based annotations to assist with the cataloguing of 3D cultural heritage artefacts.

7 Conclusion

In this paper, we have provided an overview of 3DSA – a Web-based semantic annotation system for 3D digital surrogates in the cultural heritage domain – and described the solution we have adopted to assist museum curators in cataloguing Greek vases. 3DSA enables the crowd-sourcing of semantic annotations, using the Greek Vases Ontology. Furthermore, it provides a simple interface by which museum experts can define domain-specific rules, which are then used to perform probabilistic reasoning. The experimental results of our classification mechanism are particularly encouraging and demonstrate the increased efficiency and accuracy of our approach compared to classical rule-based reasoning.

Future work will focus on two particular challenges. Firstly, we will investigate novel ways of encouraging lay users to take part in the collective annotation of 3D cultural heritage artefacts. At the same time, we will incorporate active

learning techniques to detect and improve the correctness of the crowd-sourced annotations. Secondly, we intend to deploy the 3DSA system in other domains and evaluate its inference capabilities in contexts that are not as well defined as the current application use case.

Acknowledgments. The work described in this paper has been funded by Microsoft Research Asia through the e-Heritage program and through a University of Queensland Major Equipment and Infrastructure (MEI) grant. The Open Annotations Collaboration component has been funded by the Andrew W. Mellon Foundation.

References

1. Rowe, J., Razdan, A.: A Prototype Digital Library for 3D Collections: Tools to Capture, Model, Analyze, and Query Complex 3D Data. In: Proc. of Museums and the Web 2003, Charlotte, NC, US, p. 16 (2003)
2. Hunter, J., Schroeter, R., Koopman, B., Henderson, M.: Using the Semantic Grid to Build Bridges between Museums and Indigenous Communities. In: Proc. of Global Grid Forum: Semantic Grid Applications, Honolulu, Hawaii, US, pp. 46–61 (2004)
3. Koller, D., Frischer, B., Humphreys, G.: Research challenges for digital archives of 3D cultural heritage models. *J. on Computing and Cultural Heritage* 2, 1–17 (2009)
4. Attene, M., Robbiano, F., Patane, G., Mortara, M., Spagnuolo, M., Falcidieno, B.: Semantic Annotation of Digital 3D Objects. In: Proc. of 2nd Int. Conf. on Semantic and Digital Media Technologies, Genoa, Italy (2007)
5. Chun, S., Cherry, R., Hiwiller, D., Wyman, B.: Steve.museum: An Ongoing Experiment in Social Tagging, Folksonomy, and Museum. In: Proc. of Museums and the Web 2006, Albuquerque, New Mexico (2006)
6. Cena, F., Carmagnola, F., Cortassa, O., Gena, C., Wang, Y., Stash, N., Aroyo, L.: Tag interoperability in cultural web-based applications. In: Proc. of 19th ACM Conference on Hypertext and Hypermedia, Pittsburgh, PA, US, pp. 221–222 (2008)
7. Chan, S.: Tagging and Searching – Serendipity and museum collection databases. In: Proc. of Archives & Museum Informatics, San Francisco, California, pp. 87–99 (2007)
8. Oomen, J., Aroyo, L.: Crowdsourcing in the cultural heritage domain: opportunities and challenges. In: Proc. of the 5th Int. Conf. on Communities and Technologies, pp. 138–149 (2011)
9. Yu, C.-H., Groza, T., Hunter, J.: High Speed Capture, Retrieval and Rendering of Segment-Based Annotations on 3D Museum Objects. In: Airong, J. (ed.) ICADL 2011. LNCS, vol. 7008, pp. 5–15. Springer, Heidelberg (2011)
10. Domingos, P., Kok, S., Lowd, D., Poon, H., Richardson, M., Singla, P.: Markov logic. In: De Raedt, L., Frasconi, P., Kersting, K., Muggleton, S.H. (eds.) Probabilistic ILP 2007. LNCS (LNAI), vol. 4911, pp. 92–117. Springer, Heidelberg (2008)
11. Garrette, D., Erk, K., Mooney, R.: Integrating Logical Representations with Probabilistic Information using Markov Logic. In: Proc. of the Int. Conf. on Computational Semantics, Oxford, England, pp. 105–114 (2011)

12. Albertoni, R., Papaleo, L., Pitikakis, M., Robbiano, F., Spagnuolo, M., Vasilakis, G.: Ontology-Based Searching Framework for Digital Shapes. In: Meersman, R., Tari, Z. (eds.) OTM-WS 2005. LNCS, vol. 3762, pp. 896–905. Springer, Heidelberg (2005)
13. Vavalis, M., Pitikakis, M.: Archaeology and cultural heritage awg focus k3d (2009), <http://www.slashdocs.com/ikxupp/focus-k3d-awg-archaeology-and-cultural-heritage.html>
14. Doerr, M., Tzompanaki, K., Theodoridou, M., Georgis, C., Axaridou, A., Havemann, S.: A repository for 3D model production and interpretation in culture and beyond. In: Proc. of the 11th Int. Conf. on Virtual Reality, Archaeology and Cultural Heritage, Paris, France, pp. 97–104 (2010)
15. Karmacharya, A., et al.: ArchaeoKM: Managing Archaeological data through Archaeological Knowledge. In: Proc. of the Conference on Computer Applications and Quantitative Methods in Archaeology, Granada, Spain (2010)
16. Kauppinen, T., Puputti, K., Paakkarinen, P., Kuitinen, H., Vaatainen, J., Hyvonen, E.: Learning and Visualizing Cultural Heritage Connections between Places on the Semantic Web. In: Proc. of the Workshop on Inductive Reasoning and Machine Learning on the Semantic Web, Heraklion, Greece (2009)
17. Damova, M., Dannells, D.: Reason-able View of Linked Data for Cultural Heritage. In: Dicheva, D., Markov, Z., Stefanova, E. (eds.) Software, Services and Semantic Technologies. AISC, vol. 101, pp. 17–24. Springer, Heidelberg (2011)
18. Wielemaker, J., Hildebrand, M., van Ossenbruggen, J., Schreiber, G.: Thesaurus-Based Search in Large Heterogeneous Collections. In: Sheth, A.P., Staab, S., Dean, M., Paolucci, M., Maynard, D., Finin, T., Thirunarayan, K. (eds.) ISWC 2008. LNCS, vol. 5318, pp. 695–708. Springer, Heidelberg (2008)
19. Schreiber, G., et al.: Semantic annotation and search of cultural-heritage collections: The MultimediaN E-Culture demonstrator. *Journal of Web Semantics* 6(4), 243–249 (2008)
20. Bucher, B., Falquet, G., Clementini, E., Sester, M.: Spatial relations and properties for semantically enhanced 3D city models. In: Proc. of the Workshop on 3D Issues in Urban and Environmental Systems, Madrid, Spain (2012)
21. Hmida, H.B., Cruz, C., Boochs, F., Nicolle, C.: Knowledge Base Approach for 3D Objects Detection in Point Clouds Using 3D Processing and Specialists Knowledge. *Int. J. on Advances in Intelligent Systems* 5(14), 1–14 (2012)
22. Kapahnke, P., Liedtke, P., Nesbigall, S., Warwas, S., Klusch, M.: ISReal: An Open Platform for Semantic-Based 3D Simulations in the 3D Internet. In: Patel-Schneider, P.F., Pan, Y., Hitzler, P., Mika, P., Zhang, L., Pan, J.Z., Horrocks, I., Glimm, B. (eds.) ISWC 2010, Part II. LNCS, vol. 6497, pp. 161–176. Springer, Heidelberg (2010)
23. Kasneci, G., Ramanath, M., Sozio, M., Suchanek, F.M., Weikum, G.: STAR: Steiner-Tree Approximation in Relationship Graphs. In: Proc. of the IEEE Int. Conf. on Data Engineering, Washington, DC, US, pp. 868–879 (2009)
24. Wang, X.-Y., Lv, T.-Y., Wang, S.-S., Wang, Z.-X.: An ontology and SWRL based 3D model retrieval system. In: Li, H., Liu, T., Ma, W.-Y., Sakai, T., Wong, K.-F., Zhou, G. (eds.) AIRS 2008. LNCS, vol. 4993, pp. 335–344. Springer, Heidelberg (2008)
25. Rodriguez-Echavarria, K., Morris, D., Arnold, D.: Web based presentation of semantically tagged 3D content for public sculptures and monuments in the UK. In: Proc. of the 14th Int. Conf. on 3D Web Technology, Darmstadt, Germany, pp. 119–126 (2009)
26. Serna, S.P., Schmedt, H., Ritz, M., Stork, A.: Interactive Semantic Enrichment of 3D Cultural Heritage Collections. In: Proc. of the 13th Int. Symp. on Virtual Reality, Archaeology and Cultural Heritage, Brighton, UK, pp. 33–40 (2012)