The Knowledge Engineering Review, Vol. 27:1, 33–55. © Cambridge University Press, 2012
doi:10.1017/S026988891100021X

Semantically interoperable three-dimensional scientific objects

MARIOS PITIKAKIS[1,2], MANOLIS VAVALIS[1,2] and CATHERINE HOUSTIS[1,2]

[1]*Department of Computer and Communication Engineering, University of Thessaly, GR 38221 Volos, Greece;*
[2]*Laboratory for Information Systems and Services, Centre for Research and Technology—Thessaly, GR 38500 Volos, Greece;*
e-mail: mpitikak@inf.uth.gr, mav@inf.uth.gr, houstis@inf.uth.gr

Abstract

In recent years, digital content of three-dimensional (3D) scientific objects has become widespread and is made available in a plethora of on-line scientific repositories. A systematic and formal approach becomes necessary to represent the knowledge/information related to these objects, in order to facilitate their search, retrieval and reuse. In this paper, we employ semantic interoperability to manage 3D scientific object content in a large-scale framework of scientific applications that consists of 3D shape models, associated tools and resources. We use an advanced ontological organization for their metadata, powerful reasoning engines for their search and retrieval and we elucidate several crucial issues in the design and implementation of 3D knowledge-based management systems.

1 Introduction

Three-dimensional (3D) content is widely recognized as the upcoming wave of digital media and the success of 3D environments and applications reveal that there is a shift in the way people see and navigate the Internet. Users are becoming more and more actively involved in content creation and they are continuously demanding tools that are intuitive and effective to create, share, retrieve and reuse 3D content. Moreover, there is a wide range of applications spanning from entertainment and education to scientific visualization that use 3D content and the decreasing cost of producing and/or collecting 3D data greatly impacts on a number of industrial and scientific sectors as well as everyday life.

Rapid technological evolution and at the same time emerging needs raise new challenges, in a variety of communities, associated with the management and access of the amounts of information carried by 3D content. In addition, knowledge and information reuse are vital components in order to maintain the competitive edge and ensure scientific or business success. 3D models are being used increasingly as a basis of new activities and significant time is spent searching for the appropriate 3D information. To address the emergence of these newly created demands, the term *semantic multimedia* have been proposed as the new paradigm that encapsulates the convergence of multimedia and knowledge technologies and is also regarded as the evolution of traditional multimedia. Semantic multimedia can make it possible to efficiently use and reuse, share and access digital content in distributed and networked environments.

Repositories of 3D content store information of different sources and contexts. 3D data do not only represent graphical aspects, they also hold a high knowledge value, either due to the expertise needed in their design or to the information content carried. This knowledge is of different kinds. Knowledge related to the geometrical and visual aspects that can be captured by a set of geometric and graphical data representing the digital object. Knowledge related to the purpose/role of the

https://doi.org/10.1017/S026988891100021X Published online by Cambridge University Press

object represented, which defines its category or functionality. Knowledge related to the application domain and the way the 3D data are represented, processed and interpreted. As pointed out by Hendler (2003), researchers may need to find out and explore results at different levels of granularity, from another part of the field or from a completely different scientific field. Although scientists are relying on the web to share their own scientific resources (tools, papers, data resources), current web technology is clearly insufficient to support collaborative e-science and a more elaborate framework is needed. Furthermore, internet-based services are shaping our economy and a new world of internet-enabled services on top of the internet of things is fast approaching. 3D models will sooner or later become the dominant thing in this world. Therefore, 3D models should not be seen as isolated entities, but as fundamental components of the future internet, at the same time, dominating the web.

In this context, the AIM@SHAPE project (Falcidieno *et al.*, 2004; AIM@SHAPE, 2005) introduced knowledge-management techniques for visual content with the aim of making explicit and shareable the knowledge embedded in multi-dimensional media, with focus on 3D content. This required the development of both ontologies and knowledge bases capable of describing 3D objects and processes, and data structures and tools used to associate semantics to 3D models. It was also necessary to build a common framework for managing, storing, reasoning, searching and interacting with the semantic content related to the knowledge domain. This framework is the Digital Shape Workbench (DSW; AIM@SHAPE, 2006a), one of the main outcomes of AIM@SHAPE, a common infrastructure which incorporates software tools, visual media databases and a digital library, all built on the basis of suitable ontologies and metadata.

This paper aims at presenting this knowledge-based approach for 3D resource management, based on Semantic Web technologies in order to construct semantically enriched components that can be integrated to form an environment for collaborative services. This provides a unified infrastructure that is necessary for efficiently searching and retrieving 3D resources.

The rest of this paper is organized as follows: Section 2 contains the motivation behind our efforts and sets the appropriate background. Section 3 reviews related work. Section 4 shortly describes the conceptualization of the domain by using suitable ontologies. Section 5 provides an overview of the DSWs main components. Section 6 describes, in detail, the design, implementation, integration and general technological issues associated with our knowledge infrastructure framework. Section 7 briefly comments on the evaluation of the developed system as this was deployed and still operates. Finally, in Section 8 we present our concluding remarks, the proposed future work and the envisaged research road-mapping toward semantic interoperability for 3D content.

2 Background and motivation

3D media are digital representations of either physically existing objects or virtual objects that can be processed by computer applications. Besides their impact on entertainment and 3D web, the ease of producing and/or collecting data in digital form has caused a gradual shift of paradigm in various applied and scientific fields: from physical prototypes and experience to virtual prototypes and simulation. This shift has an enormous economic impact on a number of industrial and scientific sectors, such as Design and Manufacturing, Gaming and Simulation, Cultural Heritage and Archeology, Medical Applications and Bioinformatics, where 3D media are essential knowledge carriers.

Thanks to the technological advances, we have plenty of tools for visualizing, streaming and interacting with 3D objects, even in non-specialized web contexts. Conversely, tools for coding, extracting and sharing the semantic content of 3D media are still far from being satisfactory. Automatic classification of 3D databases, automatic 3D content annotation, and content-based retrieval has raised many new research prospects. At the same time, knowledge technologies, such as structured metadata, ontologies and reasoners, have proven to be extremely useful to support a stable and standardized approach to content sharing, and the development of these techniques for 3D content and knowledge intensive scenarios is still at its infancy.

Creating a 3D model takes much time and effort and it would therefore be greatly advantageous to be able to reuse and adapt existing 3D models instead of creating every time new ones

https://doi.org/10.1017/S026988891100021X Published online by Cambridge University Press

from scratch. Embedding semantics and domain knowledge in 3D-intensive applications can highly improve the content creation pipeline, in terms of speeding up the process and allowing a concrete reuse of valuable resources (e.g. existing content, processing tools, workflows, etc.). This has motivated the development of 3D shape retrieval mechanisms that are based mainly on 3D shape matching. That is, given a user-defined shape object, the search mechanism is able to retrieve shapes from the storage database that are considered to be similar enough, based on some criteria specific to the algorithm that performs the matching, to the object provided. These content-based 3D shape retrieval methods usually perform much better than simple text-based solutions (Kazhdan *et al*., 2004). This is because unlike text documents 3D models are much more complex objects and are not easily retrieved. In some cases, where the object has a well-defined meaning, a text-based search can be very effective. In general though, conventional text-based queries are too limited and ambiguous to be used for retrieving 3D objects.

Even content-based methods, however, are not sufficient and are characterized by relatively low precision. The problem is that these methods are mainly based on geometrical matching and provide no consideration regarding the semantics of the object to be retrieved (Min, 2004).

Although, the necessity for moving toward a representation of multi-dimensional objects which also includes semantics is easy to understand, there is no concrete methodology to use in order to achieve this. The reason is that it is a relatively new research area and the complexity involved in dealing with this type of objects is high.

When dealing with the semantic representation of shapes we face the initial problem that the knowledge carried by digital shapes is quite limited. Furthermore, digital shapes are coded through models and different data formats that have no relation to data interpretation. Capturing the knowledge that a shape encompasses means to be able to refer to a shape with regards to its meaning and its usage, without taking into account its geometric representation. However, it is important that this semantic mark-up is formalized in order to not only be understood by humans, but also to be machine-understandable. This is crucial in realizing the vision of developing intelligent agents and programs able to interoperate and access knowledge bases, dealing with multi-dimensional objects in the same way as with any other type of information in the Semantic Web today.

Another part of the problem we are faced with is how to automate this knowledge extraction from shapes. Manual annotation of semantics is tedious work in general but in the case of such complex objects it is highly inconvenient due to the mass of information and the complexity involved to represent this information in a formal manner. In order to facilitate automatic annotation we need to define rules for associating semantics to lower-level knowledge in shapes. We can use ontology structures to define these rules but we also need tools to actually perform the semantic annotation.

Semantic description of multimedia items has been mainly developed for audio, video and still images (see e.g. the activities of the W3C Multimedia Semantics Incubator). Domain-specific ontologies are focused on describing the content and the parts of a multimedia scenario, such as elements in a scene, colurs, motion duration, etc. These descriptions are defined in order to be able to categorize, retrieve and reuse multimedia elements.

Examples of domain-specific ontologies and metadata have been developed for a wide set of applications, from Cultural Heritage (Doulaverakis *et al*., 2005) to Biomedicine (Catton *et al*., 2005). FOCUS K3D (FOCUS K3D, 2008) has identified four application areas that are both consolidated in the massive use of 3D digital resources (Medicine & Bioinformatics and CAD/CAE (computer-aided design/computer-aided engineering) & Virtual Product Modeling) and emerging (Gaming & Simulation and Archeology & Cultural Heritage) for promoting the use of semantics in 3D content. Although knowledge technologies and Semantic Web technologies are well established and used for other kinds of content (e.g. text, images), no particular methods/standards are used for managing and organizing 3D content. The approaches toward domain-specific ontologies developed so far are either too general for being immediately useful in real application domains or too specific to a single domain. Therefore, it is necessary to promote and demonstrate the accessibility and viability of research outcomes in the representation and processing of semantics in 3D media in a number of applied sectors. A fast evolution of the research in

https://doi.org/10.1017/S026988891100021X Published online by Cambridge University Press

the field is conditioned by the ability to create interdisciplinary research teams able to communicate with the application domains that produce and use 3D content. 3D is spreading out of the traditional circles of professional users and it will soon reach new audiences.

Management and retrieval of 3D content is emerging, besides AIM@SHAPE (2005) project, which we shall consider in some detail in the subsequent sections, in several other European Union projects. SALERO (2006), for example, approaches the problem of associating the semantics to 3D content but with specific focus on building an expert system for processing and re-purposing digital content for gaming applications, and 3D is one type of content considered. Since gaming is the only application field considered, the range of issues to be addressed for 3D content and its semantic descriptions are limited to a single and rather simple type of 3D representation. VICTORY (2006) targets 3D content from the point of view of the development of a search engine to retrieve 3D and associated multimedia-distributed content in peer-to-peer and mobile networks. In this case, however, the issue addressed is concerning more the retrieval of the content rather than the modeling of the semantic aspects of 3D objects. SCULPTEUR (2005) developed new modes of information retrieval and collection exploration by combining Semantic Web and content-analysis techniques to allow searching by concept, metadata and content. 3D-COFORM (2008) aims to establish 3D documentation as an affordable, practical and effective mechanism for long-term documentation of tangible cultural heritage. FOCUS K3D (2008) aims to foster the comprehension, adoption and use of knowledge intensive technologies for coding and sharing 3D media content in consolidated and emerging application communities.

Concluding the discussion in this section we would like to point out that semantic 3D media, as the evolution of traditional graphics media, makes it possible to use process and share 3D content of multiple forms, endowed with some kind of intelligence, in distributed or networked environments. The success of semantic 3D media largely depends on the ability of advanced systems to provide efficient and effective search capabilities, analysis mechanisms, intuitive reuse and creation facilities, concerning the content, semantics and context.

3 Related work

The advent of Semantic Web technologies represents an important advance in creating networks of knowledge, based on various on-line available resources. In practice, different methods have been adopted for documenting 3D content. There is a desperate need and a deep awareness of the necessity of harmonizing standards. As yet, interoperability is obtained by reducing the interoperable information to the Dublin Core (DC) metadata (Dublin Core, 2006). The problem, however, is that the DC was created based on different objectives, and does not preserve the richness of existing repositories, and this is even more evident in 3D digital content.

In the present context, metadata about 3D models exist in a number of different levels. Some metadata are concerned about the origins of an individual model, including the producer, the methods of capture, the conditions at the time of capture, information on the settings used for equipment, perhaps on the algorithms used, etc. There may also be items of legal interest connected to the artefact: owner, copyright status, fee for reuse, etc. Other metadata contain links and relationships to data, which are part of the same collection, recorded at the same time, in the same season, by the same person, etc. There are several approaches; however, none of them deal with the difficulties that inherently exist in capturing 3D shape knowledge.

All standardization efforts are classified along two principal dimensions, content type (i.e. the standards subject domain) and semantic depth. A 3D digital object is a collection of all the geometric, structural and visual data (model representation) as well as semantic features and properties.

The management of collections of 3D objects is at an early stage and there is significant debate whether the approaches adopted for digital libraries are suitable also for this kind of digital objects. There have been significant efforts to define metadata formats and introduce semantic information. From the perspective of documenting historic and cultural heritage artefacts in museums and archives, the CIDOC-CRM (International Council of Museums, 2009) initiative has

https://doi.org/10.1017/S026988891100021X Published online by Cambridge University Press

reached ISO standard (2007). Specifically the ISO 21127:2006 establishes guidelines for the exchange of information between cultural heritage institutions and is based on input from the International Council of Museums. This approach has been adopted by a number of other projects and was the basis of the ontological work included in the EPOCH (Excellence in Processing Open Cultural Heritage) Common Infrastructure (EPOCH, 2008).

It is worth to mention that the CIDOC-CRM is an object-oriented extensible data model and can be regarded on one side as ontology to describe the semantics behind data structures of cultural repositories, and on the other side it can be seen as a top-level ontology for integrating the various terminologies. It provides definitions and a formal structure for describing the implicit and explicit concepts and relationships used in cultural heritage documentation.

Moving Picture Experts Group (MPEG-7; ISO-IEC Moving Picture Experts Group working group, 2004) is a standard for describing multimedia content. It addresses a wide variety of media types including: still pictures, graphics, 3D models, audio, speech, video and combinations of these. Furthermore, MPEG-7 also provides an ontology (Hunter, 2001), which embodies a general and large representation of metadata. The Visual Descriptors Ontology (Bloehdorn *et al.*, 2004), written in RDFS (2009), aims to offer a more extensive description of the visual part of MPEG-7. The Core Ontology for Multimedia (Arndt *et al.*, 2007) is another ontology that extends MPEG-7 to provide richer multimedia semantics by using generic software patterns, which create a layer between MPEG-7 concepts and domain-specific interpretations.

There are efforts toward a generalized multimedia ontology (AceMedia, 2005), which represent the challenge of unifying concepts among domain-specific and top-level ontologies. Attempts of combined ontologies also exist. For example, as shown in Hunter *et al.* (2003), an ontology was developed based on the ABC ontology and model (Lagoze & Hunter, 2001), which imported the MPEG-7 ontology, the MPEG-21 Rights ontology (Pereira, 2001) and the CIDOC-CRM ontology. The metadata input was a combination of attributes coming from all the above ontologies, providing semantic interoperability across communities. Physical artefacts and digital objects were described from a museum/library/archive record point of view, with rights/access constraints, and were annotated with spoken/textual descriptions. Another kind of standard deals with data exchange formats. By standardizing data exchange formats for 3D objects, data files can be interpreted to build data structures for use in a variety of applications. Currently, X3D (2009; succeeding the earlier Virtual Reality Modeling Language) and COLLADA (2009) are two popular standards for encoding 3D objects into common formats for sharing between applications.

For modeling and visualization purposes, X3D offers many types of geometries, appearances and rigid transformations (i.e. rotation, translation and scaling). The geometries include 2D primitives (such as points, arcs, circles, ellipses, lines and polygons) and 3D primitives (such as boxes, cones, cylinders and spheres). If complex 3D geometries are desired, polygon meshes can be employed. Besides primitives for general modeling purposes, specialized components, such as Humanoid Animation (2009) and NURBS for computer-aided design (CAD) applications, exist in X3D as well.

COLLADA is increasingly used in 3D content creation while X3D has particular legacy value and familiarity in the field. COLLADA defines an Extensible Markup Language (XML)-based schema to make it easy to transport 3D data between applications—enabling diverse 3D authoring and content processing tools to be combined into a production pipeline. It is also an extensible format that is foreseen to grow to support increasingly sophisticated 3D features, as they evolve (Arnaud & Barnes, 2006).

The development of the infrastructure of the DSW, considered in Section 5, is mainly based on a knowledge management system (KMS) and a search mechanism. Sections 3.1and 3.2 summarize the related work in these two topics.

3.1 Knowledge base systems

KMSs are designed to allow users to access and utilize rich sources of data, information and knowledge stored in different forms, but also to support knowledge creation, knowledge transfer and continuous learning for the knowledge workers. Recently KMSs, unlike databases, have

https://doi.org/10.1017/S026988891100021X Published online by Cambridge University Press

aimed beyond the mere administration of electronic information; they now aim at fostering learning processes, knowledge sharing, collaboration between knowledge workers irrespective of their location, etc.

The support of querying facilities has always been a primary requirement for repositories of any kind. The proliferation of knowledge caused by the widespread use of the Web as a knowledge communication platform has posed the same and even more imperative requirements for performing queries and locating resources into the vast information space. We believe that the addition of explicit semantics can improve search. However, the data models used to represent and encode knowledge on the Web differ from the traditional data structures. Resource Description Framework (RDF), RDF Schema (RDFS) and Web Ontology Language (OWL) are the W3C standards used to encode web-based data. Thus, the functionality a querying language should support is according to the structure and peculiarities of the new paradigms. Practical Description Logic (DL) systems such as Racer (Galinski *et al.*, 2005) offer a functional Application Programming Interface (API) for querying a knowledge base.

Various Knowledge Base Systems (KBSs) have been developed for storing, reasoning and querying Semantic Web information. They differ in a number of important ways. For instance, many KBSs are memory-based while others use secondary storage in order to provide persistence. Another key difference is the degree of reasoning provided by the KBS. Some KBSs only support RDF/RDFS inference while others aim at providing OWL reasoning.

HP Labs has developed Jena, a Java framework for building Semantic Web applications (McBride, 2002; available as open source software under a Berkeley Software Distribution (BSD) license). Jena implements APIs for dealing with Semantic Web building blocks such as RDF and OWL. It contains a rich set of features for dealing with RDF including an inference engine, the ability to reason using custom rules and OWL-based ontology processing. It provides a semantic layer which abstracts a lot of the underlying knowledge that is required to work with the RDF or OWL syntax. The Jena Ontology API is language-neutral and as such it supports the common elements found in ontology languages. The inference API allows access to the available reasoners for OWL but to utilize all available functionality one should communicate directly with the reasoner.

HP Labs also maintains Joseki (W3C RDF Data Access Working Group, 2009), a web application for publishing and querying RDF models on the web, which is also available under a BSD license. It is built on Jena and, via its flexible configuration, allows a Model, represented as a set of files or within a database, to be made available on a specified Uniform Resource Locator and queried using the Protocol and RDF Query Language (SPARQL; W3C Recommendation, 2008).

3store (ATK project, 2008) is maintained by a development team at the University of Southampton. 3store is a core C library that uses My Structured Query Language to store its raw RDF data and caches. 3store is available under the General Public License, and Redland under the LGPL or Apache version 2.0 licenses. It provides access to the RDF data via RDF Data Query Language (Seaborne, 2004) or SPARQL over Hypertext Transfer protocol (HTTP), on the command line or via a C API. 3store can infer RDF and RDFS entailments and can also communicate through an Apache module. The results are cached in an in-memory store for browsing.

Sesame (Aduna Open Source project, 2009) is an open source Java framework for storing, querying and reasoning with RDF and RDFS and is distributed under the Lesser General Public License. It can be used as a database for RDF and RDFS, or as a Java library for applications that need to work with RDF internally. Sesame provides application developers a toolbox that contains the necessary tools to parse, interpret, query and store ontologies and metadata. Sesame can be used as a server with which client applications or human users can communicate over HTTP. It supports RDFS inferencing and can support OWL inferencing and querying through the OWLIM Semantic Repository (Kiryakov *et al.*, 2005). OWLIM is a high-performance semantic repository, packaged as a Storage and Inference Layer for the Sesame RDF database. Much like Jena, Sesame can use open source Relational Database Management Systems (RDBMSs) in addition to its in-memory database.

IBMs Integrated Ontology Development Toolkit (IBM Corporation, 2004) is a toolkit for ontology-driven development. It includes an Ontology Definition Metamodel (EODM) implemented

https://doi.org/10.1017/S026988891100021X Published online by Cambridge University Press

in Eclipse Modeling Framework (Steinberg *et al*., 2008) and an OWL ontology repository named Scalable Ontology Repository (SOR). EODM is the run-time library that allows the application to read and serialize an RDFS/OWL ontology in RDF/XML format, manipulate an ontology using Java objects, call an inference engine and access inference results, and transform between ontology and other models. SOR is an OWL storage, inference, and query system based on RDBMS. It supports Description Logic Program, a subset of OWL, DL and the SPARQL language.

More details about existing publicly available data store systems built specifically for the Semantic Web together with their performance evaluation can be found in ATK project (2008 and Lee (2004).

Commercial solutions for semantic data stores are also available. One of the more advanced is provided by the Enterprise Edition of ORACLE Database 11g. It provides management capabilities with native support for RDF/RDFS/OWL and inference using OWL and RDFS semantics, querying of RDF/OWL data and ontologies using SPARQL-like graph patterns embedded in Structured Query Language (SQL), and ontology-assisted querying of relational data. A Java API and SPARQL support via an Oracle Jena adaptor is also available.

Another business oriented solution is Institute of Technology and Management e-Knowledge from MONDECA, which is a knowledge representation management application based on Semantic Web technology and ontologies. It supports XML, RDF, Simple Knowledge Organization System, OWL and an optional automated reasoning module using the data in the database and a reasoning and inference rules editor.

3.2 Three-dimensional search engines

Traditional text-based document search schemes are not effective for the classification, search and reuse of 3D data (Kazhdan *et al*., 2004), while numerous research efforts focus on content-based retrieval, based on methods that measure shape similarity between 3D models (Veltkamp, 2001; Daras *et al*., 2004; Funkhouser *et al*., 2005). The only effective means to perform context-based retrieval rely on textual annotations of media (e.g. keywords), which are commonly inserted manually and constitute only a negligible portion of the information stored.

As 3D objects are steadily growing in number and their popularity is rapidly increasing, the academic community has started focusing its attention on search system at the National Institute of Multimedia Education, Japan (Chen *et al*., 2003), the 3D model similarity search engine at University of Konstanz, the 3D retrieval engine at Utrecht University, the GEORGLE system at Technion, Israel Institute of Technology (Leifman *et al*., 2005), the search engine for CAD models at Purdue University (Jiantao & Ramani, 2007) and the AIM@SHAPE geometric search engine (GSE). Finally, we mention the 3D search engine developed by the ongoing project VICTORY (2006). VICTORY aims to develop a distributed 3D search engine that will enable searching, accessing and retrieving 3D and other types of digital content in a distributed object repository, through Peer-to-Peer networks. The 3D search demonstrated in the project web site makes use of five different global descriptors, which can be used as different alternatives in the search. The search, however, does not allow for combination, weighting or other kinds of user interactions with the search engine.

Very few commercial 3D search engines are currently available. As an example we mention the Geolus Search engine, commercialized by the Siemens PLM Software, represents the current state-of-the-art in terms of commercial search engines for PLM.

However, geometry-based retrieval is just one of the issues that have to be addressed, while the peculiarities of the domain of usage make the general problem inherently complex. The knowledge is not solely carried by digital media, but also by the means used to acquire and transform them. Every context adds in principle another semantic layer of information to shape resources.

Current 3D media repositories are focusing on the geometric aspect of models, and not on the knowledge they represent. Although on the geometric or structural representations of resources there are numerous approaches, at the semantic level very little work has been done until now. We believe that the addition of explicit semantics can further improve search results.

https://doi.org/10.1017/S026988891100021X Published online by Cambridge University Press

In the last few years there has been a considerable increase of interest for techniques to extract and stream knowledge embedded into multimedia content, ranging from basic research efforts to projects seeking an integrated effort at European level (e.g. SCHEMA, 2002; AceMedia, 2005; FOCUS K3D, 2008).

4 Ontology development for three-dimensional models

The use of 3D data is not only related to visual aspects and rendering processes, but it involves also an adequate representation of domain knowledge to exploit and preserve both the expertise used for their creation and the information content carried.

3D semantic modeling includes activities like the acquisition and reconstruction of objects, providing high geometric accuracy in the digital models, extraction of features and/or model properties, portability and flexibility in applications, while minimizing human interaction during the modeling process. The association with semantics is also crucial to visualize the models properly and retrieve them efficiently from large repositories. Efficient retrieval implies equipping 3D content with metadata related to both the whole object and its sub-parts, developing automatic metadata extraction tools and shape similarity mechanisms to compare objects, providing best practices assisting the processing phase.

The development of the AIM@SHAPE DSW (AIM@SHAPE, 2006b) required the conceptualization of the domains and the precise characterization of the resources. In this sense, the AIM@SHAPE effort can be seen as a step toward contributing to the goals of the Semantic Web as the mean to share scientific resources.

The complexity and the extent of the 3D domain makes unsuitable the use of a single monolithic ontology and it leads the way in building a framework where different ontologies are adopted to represent different facets of specific domain applications and usage scenarios.

One of our objectives was to formalize the domain knowledge into context-dependent ontologies and introduce knowledge-management techniques in 3D shape modeling, with the aim of making explicit and shareable the knowledge embedded in digital models.

The formalization of knowledge related to the geometric and structural representation of digital models is the basic step in building application-specific ontologies. The geometry and structure of an object always remains the same while its descriptions can differ according to the various contexts, and it should be captured by a set of metadata that fully describes the properties of the representation used.

The framework introduced by AIM@SHAPE relies on the conceptualization of different aspects of an object. Ontology-driven metadata provide an expressive characterization of shapes, providing different levels of representation of a resource. To be more precise, a shape can be described as a simple resource, and/or by its geometry, its structure, its semantics, depending on the application domain.

In AIM@SHAPE we have specified a high level ontology for digital shapes, the *Common Shape Ontology* (CSO; Vasilakis *et al.*, 2010) and whose role is to express in a formalized way the knowledge about digital shapes that is common to all domain-specific scenarios used in the project. It targets different kinds of multimedia content, ranging from 2D/3D images to videos, 3D models and 3D animations, and maintains top-level information that is sharable and usable in different domains. The CSO deals with 3D models as a key resource type and it has been designed and used for a full characterization of shapes in the AIM@SHAPE Shape Repository (SR; AIM@SHAPE, 2006b). An overview of the CSO high level structure, where the most important concepts are shown, is given in Figure 1.

This ontology is detailed enough to be used and instantiated directly, but also general enough to be referred to and extended by other specialized domain ontologies (hence the name common). In fact it constitutes the foundation for the three domain-specific ontologies also developed in the AIM@SHAPE, namely the *Product Design Ontology* (PDO; Catalano *et al.*, 2008), the *Virtual Human Ontology* (VHO; Gutirrez *et al.*, 2007) and the *Shape Acquisition and Processing Ontology* (SAPO; Albertoni *et al.*, 2006). Our aim was to address multiple contexts and applications in a semantic-aware level of representation where the shape knowledge can be exploited.

https://doi.org/10.1017/S026988891100021X Published online by Cambridge University Press

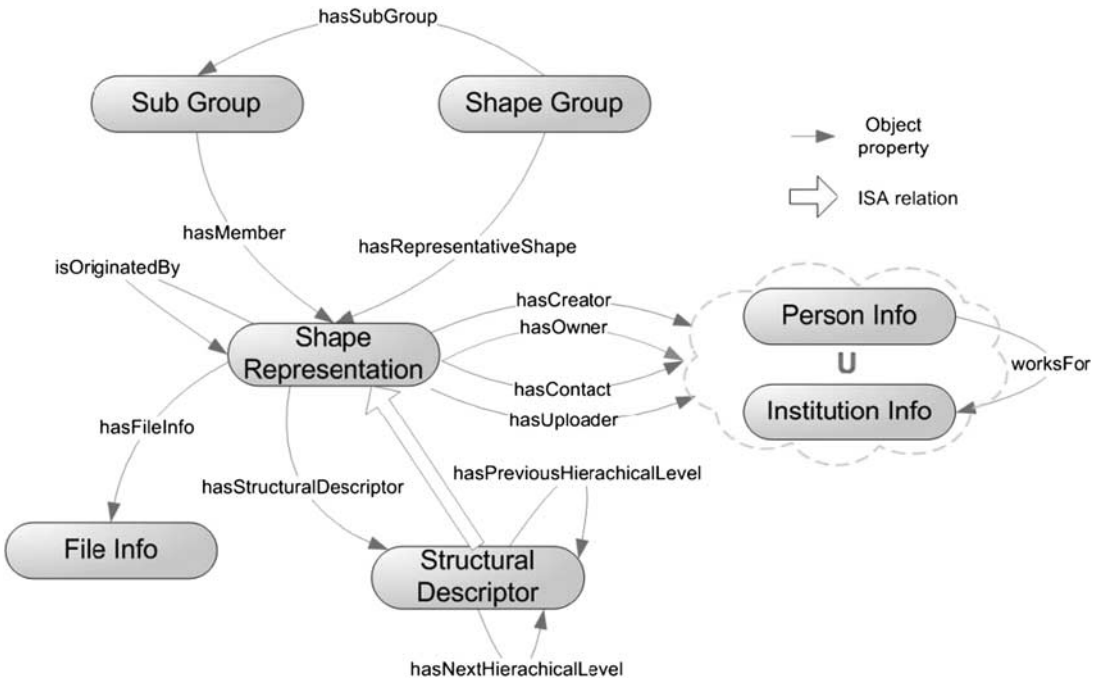


Figure 1 An overview of the high-level structure of the Common Shape Ontology

The VHO aims at organizing the knowledge and data related to research and applications in the field of virtual environments and humans: the modeling and analysis of virtual human body and their animation and interaction with virtual objects.

The SAPO intends to formalize the knowledge pertaining to the development, usage and sharing of hardware and software tools and shape data by researchers and experts in the field of shape acquisition and processing.

The objective of the PDO is to guide researchers and experts in the development of tools and methods for: supporting industrial product design and engineering analysis, dealing with knowledge concerning shape processing methods and algorithms and knowledge about processes and workflows regarding product development phases.

Concepts that are shared by all domain ontologies have lead also to the creation of a Common Tool Ontology (CTO) related to shape processing tools.

A representation using ontologies is able to provide an expressive characterization of shapes at different levels of abstraction and ensures that existing tools, such as DL reasoners (Baader *et al*., 2003), can be used to reason on the repository and deduce information that is implicit in a digital model.

The OWL (W3C Recommendation, 2004) was used for the representation of shape knowledge in all the developed ontologies. OWL has been chosen as it currently represents a standard for designing ontologies and had enough expressive power for our needs.

All the ontologies in the AIM@SHAPE Network of Excellence were developed/modeled with the help of domain experts and project partners. It was an iterative and time consuming process that took several months to reach the desired consensus among the experts.

5 The digital shape workbench: an overview

One of the main achievements of the AIM@SHAPE project is the implementation of the DSW where resources and knowledge are integrated into a unified interface. The main services supported by the DSW are:

(a) uploading of models, tools and bibliographic references;
(b) advanced searching, browsing and downloading of resources;

https://doi.org/10.1017/S026988891100021X Published online by Cambridge University Press

(c) management of resource metadata; and
(d) management and maintenance of ontologies.

The DSW consists of the data repositories (*SR* and *Tool Repository*), the KMS (*Ontology and Metadata Repository, OMR*) and a number of different ways of discovering, searching and browsing resources (*Semantic Search Engine (SSE)*, *GSE*, *Digital Library* of publications).

The SR is an open repository to the research community, populated with a collection of digital models. Its primary goal is to include a variety of standard test cases and benchmarks, enabling efficient prototyping as well as practical evaluation on real-world or large-scale models. The SR has some unique features that differentiate it from other repositories. It contains high quality models categorized and documented through ontology-driven metadata. When a model is downloaded, its thumbnails and metadata are automatically bundled into a package, which is then sent to the user on agreement to the applicable licenses. It is also possible to choose the quality of the downloaded model.

Another unique feature is the support of group models. It sometimes becomes necessary to have different versions of the same model with slightly altered properties, for example, different re-meshing algorithms applied on the same model. Models logically related to each other can be stored as a single group with a representative model at the topmost level. Other models can be stored in lower levels and each level is accompanied by a level description that gives an idea of how models in that level are related to each other. Similarly, each group is accompanied by a group description. With constant addition of models the repository has rapidly grown and it now contains more than 1000 models.

The *Tool Repository* is an inventory of software tools that can be used in different stages of digital shape processing. It mainly contains existing tools developed by the AIM@SHAPE project partners. It also combines, adapts and enhances tools developed by different research teams. The tools in the Tool Repository are organized based on a functionality hierarchy. For every tool a brief specification of its usage, limits and capabilities are provided as metadata. Currently, there are more than 80 software tools and libraries in the repository.

The *Digital Library* aims at the creation of a common repository of scientific references and technical reports, which integrate the bibliographies of the participating Institutes and will become a reference point for publications on the topic of shape modeling and semantics. Each item in the *Digital Library* is described by metadata. It does not store any copyrighted full text, but it provides a Uniform Resource Identifier to the resource at the publisher site. The goal of the *Digital Library* is to strive for quality; entries are carefully checked for their relevance, completeness and accuracy and there are more than 2800 references in the field.

The DSW also provides a common vocabulary/glossary of terms concerning digital shapes. It contains a selected set of over 370 terms and their definitions that are distinctive in the domain of shape modeling.

The goal of the *Search Engine* is to provide a facility aiding in the discovery of shape resources and tools, as well as references to these resources. The fundamental components of the DSW search mechanisms are the inference mechanism for semantic-based searching (Vasilakis *et al.*, 2005) and the geometry-based search mechanism.

The aim of the inference engine is not to restrict the user in searching only for shape resources (i.e. models, tools, references, etc.), but to be able to acquire other information relevant to the resource as well, using semantic criteria based on the ontology and its instances. This information may be either explicitly or implicitly described. Explicit information (metadata) includes datatype property values (i.e. specific alphanumeric values of properties), object property values (i.e. relations between classes or instances) or logical expressions/combinations of both. Implicit information can be inferred from explicit information about a resource.

To help the user in making efficient and appropriate queries we have developed a graphical user interface (UI), see Figure 2, that abstracts the query API and the underlying reasoner logic. The visual interface provides the means to search in an intuitive and straight-forward way without sacrificing the flexibility or expressiveness of the queries.

https://doi.org/10.1017/S026988891100021X Published online by Cambridge University Press

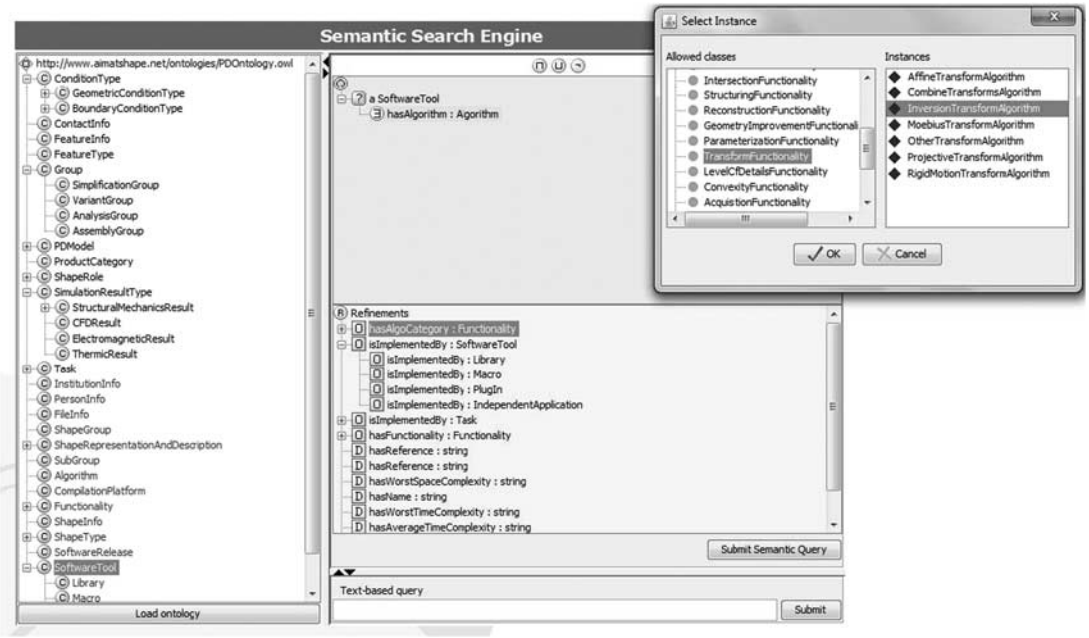


Figure 2 The Semantic Search Engine user interface

With each interaction the user is expressing some semantic criteria based on which the search mechanism will extract some information from the knowledge base. At each step the user is guided in forming these criteria by visualizing all possible options and query refinements, depending on the current context, in an intuitive way. It is also possible to explicitly store and reuse predefined user-defined queries. This is helpful in displaying the wealth of information that can be accessed and the range of interactions that are possible.

The geometry-based search mechanism is able to search the SR according to different similarity criteria and matching mechanisms (e.g. global and partial matching, sub-part correspondence, part-in-whole) and according to different perspectives (e.g. geometry-based or structure-based). Geometric similarity between shapes is evaluated as a distance between shape descriptors that capture relevant properties of the shape. Tools for computing shape descriptors and for computing distances between shape descriptors are implemented as stand-alone tools that can be plugged into the GSE. The GSE works with different matching criteria depending on the shape descriptors and the comparison methodologies available. The GSE UI is shown in Figure 3.

6 The knowledge infrastructure

The OMR constitutes the knowledge base back-end that conceptualizes and provides persistency services for the stored knowledge. As it is depicted in Figure 4, it lies in the heart of the DSW and is responsible for effectively integrating all other components. The OMR in essence is an ontology management system that allows storing, editing and accessing ontologies and ontology-driven metadata. The search mechanism's expressiveness and flexibility on performing the search is only constrained by the way that information is stored in the OMR and by the choice of the logical formalism that is used to describe this information. Hence, the OMR is closely related to the inference engine, which is necessary for semantic-based searching.

6.1 Design and implementation of the Ontology and Metadata Repository

An ontology management system is comparable to a database management system (DBMS). It allows storing and processing of ontologies and metadata via a standard interface and relieves

https://doi.org/10.1017/S026988891100021X Published online by Cambridge University Press

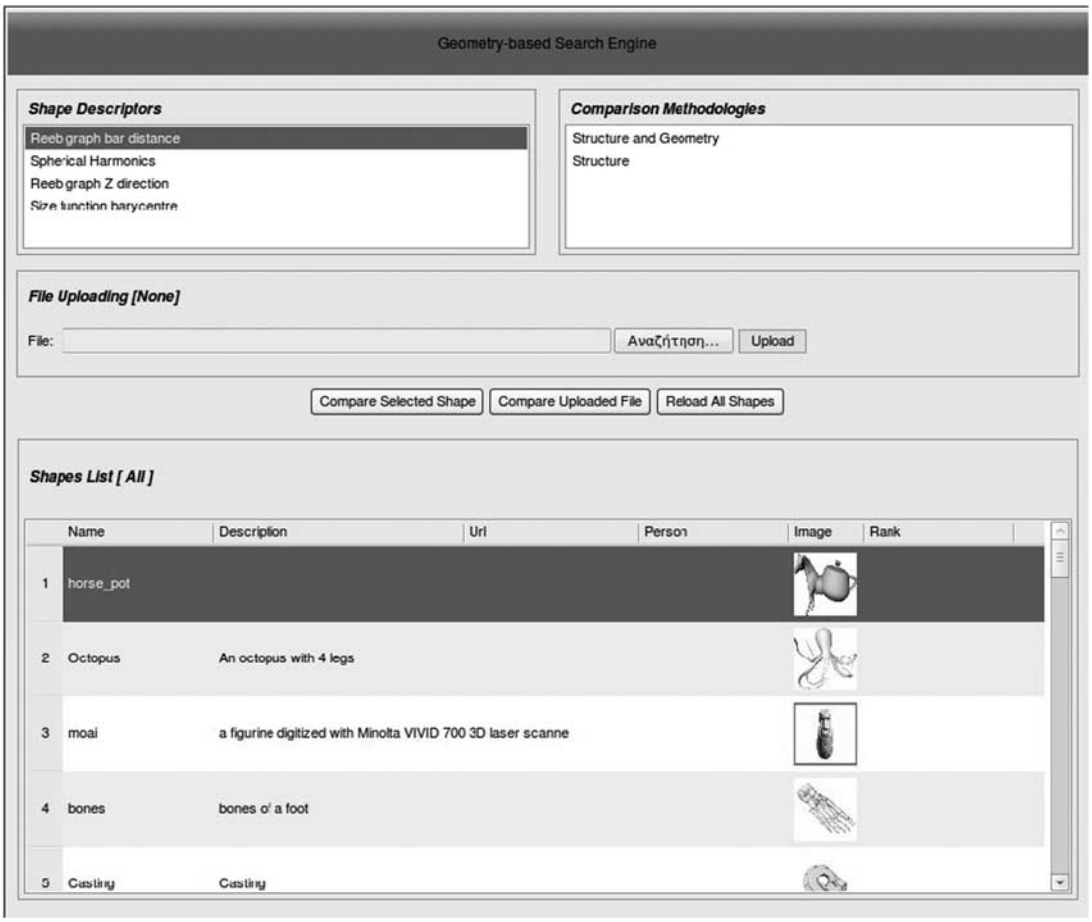


Figure 3 The Geometric Search Engine user interface

from the burden of deciding how the ontology is stored and accessed, how queries are processed and optimized, how query results are retrieved, etc. Ontology editing capabilities are not viewed as a critical component of an ontology management system since there are well-established independent ontology editors such as Protg (Noy *et al.*, 2000).

The requirements we were most interested in for such a management system were the loading of potential large ontologies (including their instances) to the storage system, the loading speed of these data into the inference engine (performance issues) and efficient browsing and viewing of both ontologies and metadata. Ontology management and evolution was also a significant concern. Transactions like inserting, updating and deleting metadata as well as altering the ontology structure were considered as vital operations.

The knowledge base was designed in a layered manner, dividing the responsibilities into various inter-dependent, but loosely coupled components. The architectural elements that comprise the software components of the OMR are identified in Figure 5.

As shown in the diagram, the core of the knowledge base comprises the OMR and the SSE (inference engine and ontology server). There are two APIs for accessing knowledge base functionality: the Query API and the Management API. The search engine interfaces with the knowledge base through the Query API while the Management module provides functionality for managing the ontology schema and metadata.

A more detailed description of the OMR architecture is shown in Figure 6. This architecture and most of its components is generic enough to be used for other kind of knowledge as well, and not only for 3D content.

https://doi.org/10.1017/S026988891100021X Published online by Cambridge University Press

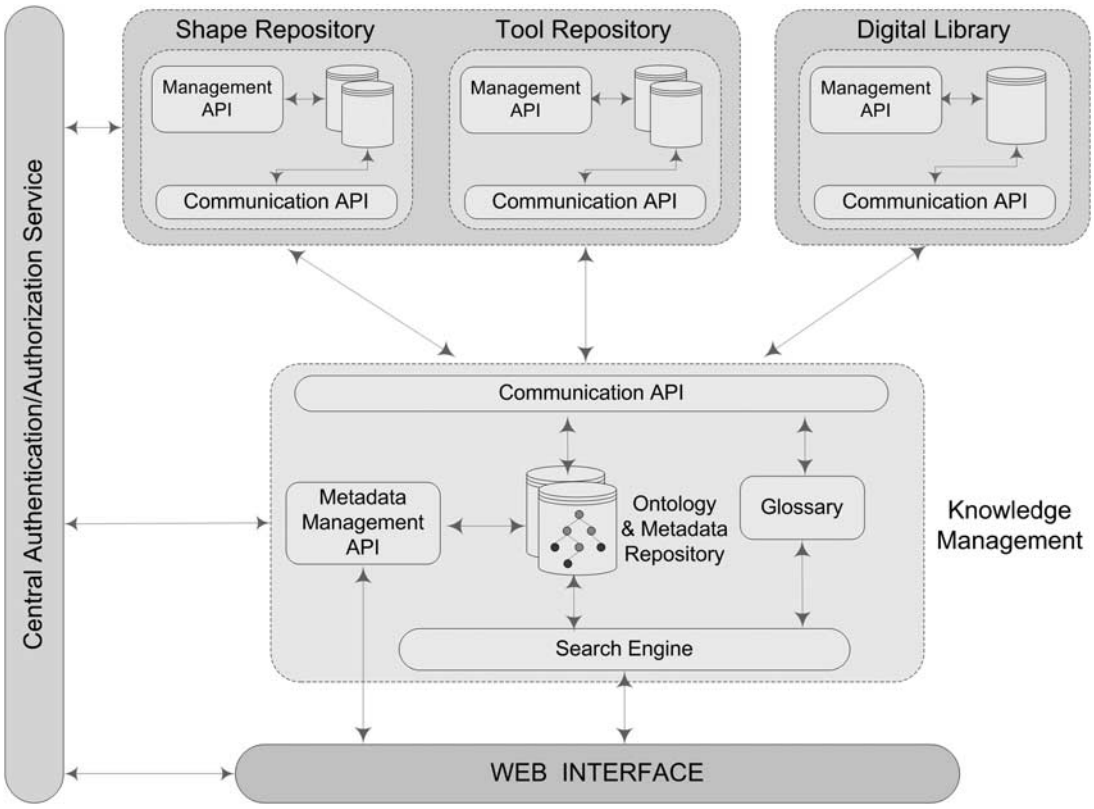


Figure 4 The overall architecture of the Digital Shape Workbench

The ontology and metadata Management Module provides the basic persistence services. More specifically, it provides a set of ontology and metadata management and maintenance activities and is capable of storing multiple ontologies.

The Management module, which is the core part of the repository, was developed on top of the Jena 2 framework (McBride, 2002) to provide management and persistence services. Jena is available as open source Java software under a BSD license and contains a rich set of features for dealing with RDF and OWL, including OWL-based ontology processing. Jena provides a semantic layer, which abstracts a lot of the underlying knowledge, that is required to work with the OWL syntax. Jena's fundamental class is the ontology model, an API for dealing with a set of RDF/OWL triples. Jena ontology models are created from OWL files and stored persistently to a RDBMS using the appropriate API. In our implementation we used Postgre Structured Query Language (PostgreSQL) as the underlying RDBMS, that is, a metadata repository for enabling the basic persistence services in a scalable and reliable fashion.

The Management module extents the Jena API and provides the following:

(a) a complete OMR replica,
(b) a database containing (among other things) part of the OMR metadata information to be used by the GSE,
(c) a database containing a copy of all the SR related metadata.

All of the above are hosted at distributed locations and in order to keep all the metadata synchronized, an API for this synchronization has been developed. The update of all the metadata repositories is synchronous, that is, all the metadata information are updated concurrently to all the repositories.

A combination of local in-memory and persistent storage approach was used. Although storing everything in-memory cannot be an effective method for manipulating extremely large volumes of

https://doi.org/10.1017/S026988891100021X Published online by Cambridge University Press

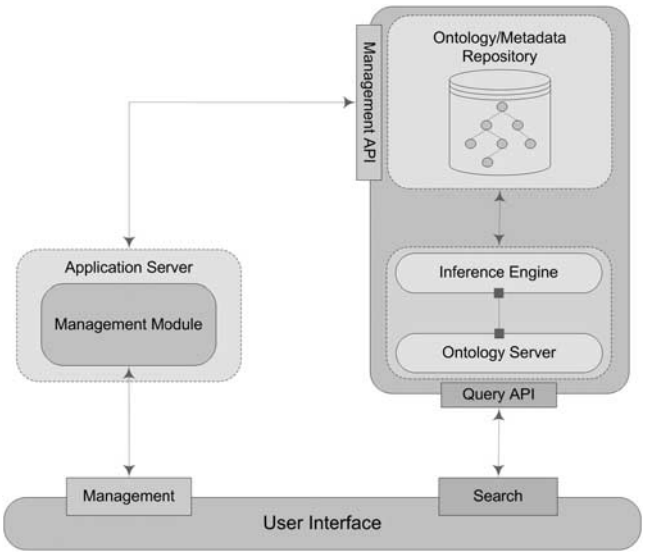


Figure 5 The Ontology and Metadata Repository general design architecture.

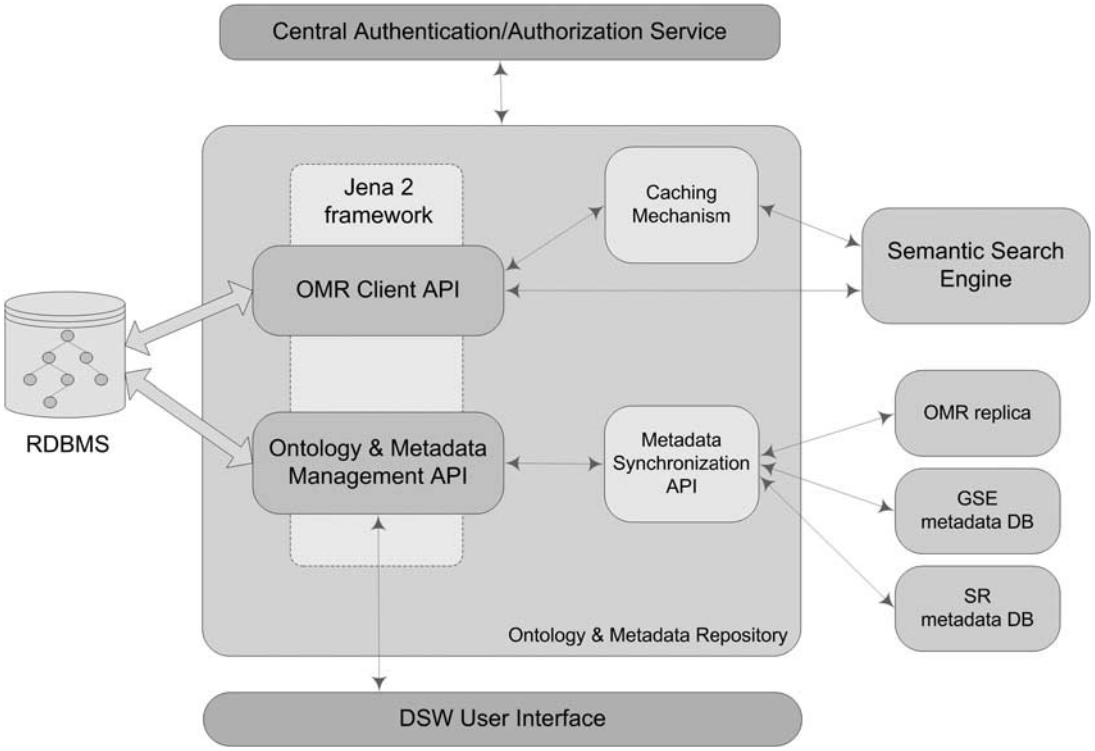


Figure 6 Detailed description of the Ontology and Metadata Repository architecture

metadata, it is very useful for caching purposes. Caching of ontology schemas is necessary and important because with imported ontologies the resulting schema that is being loaded can be quite large. Caching is performed for two distinct reasons concerning the search framework:

(a) The ontology schema that is extracted from the database is cached to minimize loading time from the persistent storage (database) to the inference engine.
(b) The representation of the schema that is being used by the UI framework. Caching the UI representation of the schema improves the necessary initialization time significantly.

https://doi.org/10.1017/S026988891100021X Published online by Cambridge University Press

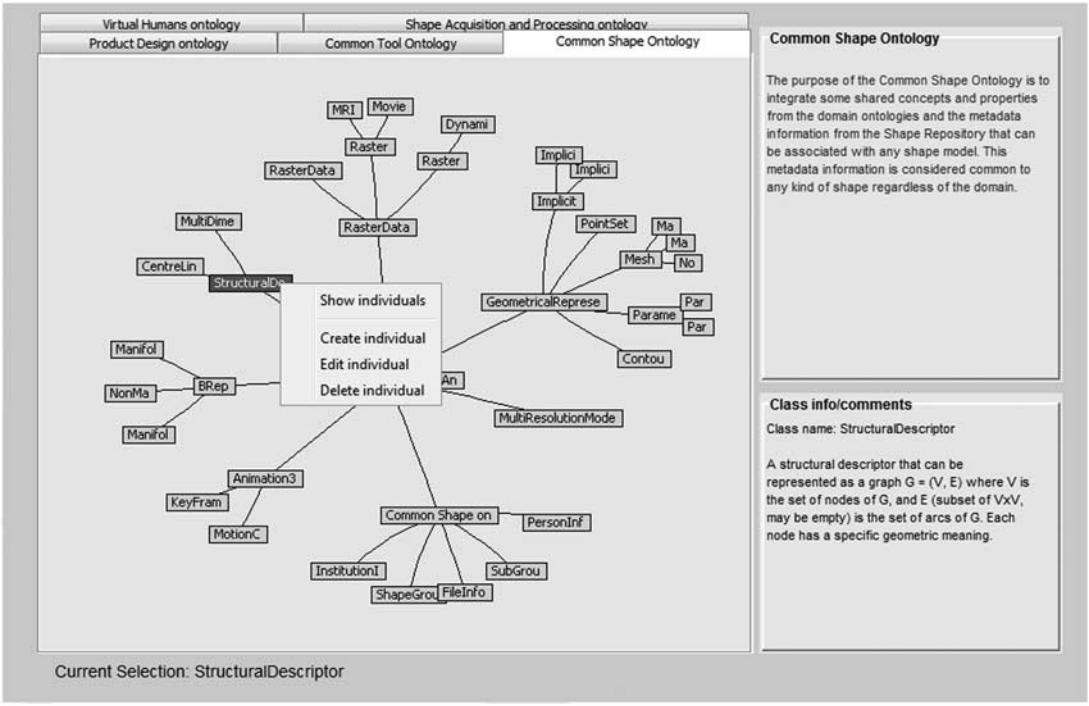


Figure 7 The hyperbolic tree visualization tool

Since ontologies were developed by different groups of people (domain ontologies and common ontologies) and continuously evolved over time, we needed a way to keep track of ontology changes. Ontology evolution and versioning is admittedly a hard research issue and out of the scope of this work. We implemented a basic set of functionalities to store and manage different versions of the same ontology (updates). Some additional information about the ontologies was also maintained like: imported ontologies (e.g. the CSO and CTO), version number, date of insertion, ontology model name. Moreover, a tracking mechanism has been implemented for logging changes to resource metadata through the DSW UI. All changes are kept for backup and security purposes.

One of the most important aspects of searching is knowing how to search. In our case, knowing how to search is related to the user's comprehension of the domain that is being conceptualized. It also depends on the user's comprehension of the way that domain knowledge has been structured according to the developed ontologies.

Our approach for displaying and navigating through ontologies is to use a hyperbolic tree visualization technique. The hyperbolic tree visualization tool was developed as a Java applet (see Figure 7) and can support management and browsing actions like displaying, inserting, editing and deleting metadata just by right clicking on a class/node in the hyperbolic tree.

Context-sensitive help has been integrated in the hyperbolic tree for any selected node, that is, the appropriate help content is displayed when a user selects a class (see bottom right panel in Figure 7). This help text is dynamically generated from the ontology schemas by reading the RDF comments that may be contained in class definitions.

Context-sensitive help has also been integrated in other OMR web pages. For example, when inserting or editing metadata, comments are displayed as floating tooltips for every property in the ontologies that has relevant RDF comments. A typical example is the allowed values of a property or the definition/meaning of a property. Furthermore, the range of a property value (e.g. integer, string, etc.) is displayed as an indication of expected values. The range is also dynamically extracted from the ontologies (i.e. the property definitions).

https://doi.org/10.1017/S026988891100021X Published online by Cambridge University Press

To sum up, the following functionalities are offered by the OMR:

(a) uploading metadata regarding shape models, tools and bibliographic references,
(b) browsing/listing all the stored ontology structures (schemas) as well as all their instances (individuals),
(c) export/download stored ontologies and metadata to OWL files,
(d) editing of resource metadata,
(e) uploading, maintenance and version tracking of OWL ontologies,
(f) interaction with metadata extraction tools for automatic metadata generation,
(g) support for pluggable reasoning modules (inference engines).

6.2 Integration with the semantic search mechanism

Jena has a set of internal/build-in reasoners (there is a micro and mini reasoner for OWL Lite and a subset of OWL Lite, respectively, and another one for OWL DL), but they are insufficient performance wise and their inference API is limited. Jena can, however, be easily combined with existing reasoner implementations and allows direct access and communication with the reasoner. In our case, the Racer Pro v1.9.1 DL reasoner (Racer Systems GmbH & Co., n.d.; Fikes *et al.*, 2003) was utilized for its inference and deductive reasoning facilities.

It is evident (see Figure 5) that the layer titled as the Query API plays a very important role. Indeed, it is the protocol that binds the SSE interface to the inference engine (i.e. Racer). Upon the specification of this protocol, and its potential for standardization, lies the flexibility and openness of the design. The benefits from choosing a protocol that is a standard, or has the potential to become one, are obvious. One could easily plug in a different inference engine that implements the same protocol and the system would still work with no required modifications. In other words, the system is truly agnostic of the actual reasoner implementation that is being used.

We have selected new Racer Query Language (nRQL) as the ontology query language for the knowledge base. The drawback to this design decision is that nRQL is not a standard and can only be used when communicating with back-ends that use some implementation of Racer. To mitigate this we have designed the Query API on top of nRQL following a layered architecture. The only element tied to nRQL is the nRQL Translator component of the API. By changing this component the API can easily be adapted to another communication protocol, for example, OWL-QL (Fikes *et al.*, 2003), much like changing a driver for a computing device would work. This ensures that the functionality provided by the Query API is not affected by changing the communication protocol and the implementation required for such a change is minimal. With this design we ensure that the functionality of the Query API is independent of reasoner specific functionality.

The integration of the SSE and the OMR is considered a critical task since the ability of efficiently answering semantic queries is based on the communication of the SSE and the OMR. For that purpose, the following integration actions took place:

(a) Development of an OMR client API for communicating with the SSE. This API provides the necessary methods for loading ontologies (schemas and metadata) from the OMR to the SSE, getting class descriptions and other ontology information, exporting ontologies (schemas and metadata) to files in order to be used by the SSE, etc.
(b) The API for metadata and ontology management was extended to support the sharing of commonly used APIs between the OMR and the SSE.
(c) The API was also modified to support caching of ontology schemas. Since performance is an important consideration for the SSE, cashing is necessary because with imported ontologies the resulting schema that is being loaded to the SSE inference engine can be fairly large. Therefore, the ontology schema is cached to minimize loading time from persistent storage (database) to the inference engine.

The Ontology Server API also supports remote management of the Ontology Server. In particular it provides support for loading specific ontologies from the OMR directly to the Inference Engine,

https://doi.org/10.1017/S026988891100021X Published online by Cambridge University Press

restarting the Ontology Server and resetting the inference engine. All the above contribute to the improved fault tolerance of the SSE.

7 System evaluation

7.1 Quantitative performance evaluation

Testing and validation of the OMR and inference engine with large-scale ontologies and metadata have been conducted. As a measure of scalability of the knowledge base we provide here an indication of the volume of metadata currently stored in the OMR. There are over 1100 shape models, more than 80 tools, more than 220 instances of Users and Institutions and over 3500 context-dependent instances in the domain and common ontologies. These also provide a solid test base for the evaluation of the efficiency of the SSE.

Our implementation choice of using a persistent storage solution (RDBMS) for storing ontologies and related metadata reduced many of the OMR performance issues to the performance of the underlying DBMS. Scalability and reliability in an RDBMS are already solved problems. The number of ontologies loaded does not affect significantly the performance of the OMR (e.g. load time, query execution) since each ontology is stored in separate tables (set of tables). To evaluate the performance of the OMR we have contacted a series of tests on a commodity desktop computer with the following characteristics: Intel Pentium 4 Processor at 3.2 GHz; 2 GB of Random Access Memory; Windows XP Professional OS; Java SDK 1.6.0 update 13; Apache Tomcat 5.5; DBMS PostgreSQL 8.1.

Table 1 shows the estimated data average loading/insertion elapsed time for the five ontologies considered above with different schemas and instance sizes. Each of the ontologies was loaded 10 times and the calculated average load time is presented in seconds with respect to its number of classes (column 2), its number of properties (column 3) and the number of instances it contains. As expected, the persistent storage (PostgreSQL) performs relatively well and displays good scalability in data loading. The results show that, as expected, the loading time is a function of both the complexity of the ontology and the number of instances.

Table 2 presents a list of three query test results and the estimated average response time measure in seconds (again the average is computed from 10 measurements). These queries represent the most common actions expected from an ontology management system or knowledge base. The increased response time concerning the last query (find and display the ontology class hierarchy) is mainly caused by the imported ontologies. The CSO and CTO import the Integrated Catalog Ontology (ICO), the Virtual Human Ontology imports the ICO and CSO and the PDO imports the ICO, CSO and CTO. In addition, the PDO extends some classes of the CSO. Therefore, the response time of such a complex class hierarchy is expected to be significantly higher.

The DSW has been running for the past 4 years while the SR and the OMR continue to increase in size. Until now we had more than 1 800 000 SR visitors and around 85 000 3D model downloads. All of the above were actively supported by the OMR. Finally, it is worth to mention

Table 1 Performance data for OMR load time with respect to ontology parameters

Data set	Ontology (no. of classes)	Information (no. of properties)	Number of instances	Average load time (s)
Info Common ontology (ICO)	2	9	289	4.95
Product Design ontology (PDO)	75	72	130	5.97
Virtual Humans ontology (VHO)	53	73	186	6.07
Common Tool ontology (CTO)	70	15	769	15.43
Common Shape ontology (CSO)	35	153	2929	105.50

OMR = Ontology and Metadata Repository.

https://doi.org/10.1017/S026988891100021X Published online by Cambridge University Press

Table 2 Performance data for OMR queries

Action	Average response time (ms)				
	ICO	CSO	CTO	VHO	PDO
Find and display all instances of a class	0.150	0.234	0.164	0.220	0.231
Retrieve an instance and display all property values	0.255	11.396	3.226	10.617	17.012
Find and display the ontology class hierarchy	0.317	3.798	6.320	8.845	21.306

OMR = Ontology and Metadata Repository; ICO = Info Common ontology; CSO = Common Shape ontology; CTO = Common Tool ontology; VHO = Virtual Humans ontology; PDO = Product Design ontology.

that all times reported in this section are in accordance with recent evaluation studies (Babik & Hluchy, 2008).

7.2 Qualitative evaluation

During the FOCUS K3D project, we contacted a questionnaire survey with the aim of collecting information on the methods and practices for 3D content use and sharing in various domains. Around 80 people answered our questionnaires from 19 different countries coming both from Industry and Education/Research. Part of the goal of the questionnaires were to identify the use of current methodologies for handling 3D content and coding knowledge related to 3D digital content in the different contexts and also to identify possible gaps and new areas of research. We summarize here some of our findings that also validated our efforts in developing the DSW.

Although the creation and/or capture of 3D models does not appear to trouble professionals (users, developers, scientists, creators of 3D content, publishers/dealers of 3D repositories, etc.), however, the management of 3D collections is at an early stage or non-existent at all. Our survey results confirmed the current situation and revealed some specific problems and shortcomings related to 3D knowledge management.

Commonly, people deal with geometric and structural aspects of 3D models while the semantic aspect appears only for the management of more complex models. Almost everybody is aware of knowledge technologies, even if half of the people questioned claim not to be familiar with them. Concerning the adopted knowledge technologies, there is a wide preference for databases and metadata and some of them use quite often taxonomies and ontologies, but not for 3D content.

Obviously the answers vary widely from scientists, researchers, developers, designers, project managers, to publishers and dealers of 3D content. Also concerning the actual amount of data, there is a huge variety ranging from just a few models (e.g. 10) to very large collections (e.g. thousands of models). Most of the data are stored on file servers or proprietary repositories. Often rather primitive approaches are used for handling 3D content (e.g. using file and folder names to encode information about the contained data) and that there is an apparent lack of information about knowledge technologies or at least a common misunderstanding or misuse of the related terminology. Tools that are specifically designed for organizing, browsing and searching 3D content are commonly not used.

People from the Archeology, Cultural Heritage and Medicine domains claim to be dissatisfied with the way they manage and store 3D content, while CAD/CAE & Virtual Product Modeling and Gaming & Simulation people are quite satisfied but they would expect an improvement. Areas mentioned for improvement or for which they feel that the use of 3D knowledge technologies can provide a solution include functionalities related to the documentation and identification of objects and parts of them, automatic extraction of geometric and semantic information from models, better visualization and improved search using semantic and/or geometric criteria.

A noteworthy observation is that 3D collections are becoming more and more demanding in terms of management, preservation and delivery mechanisms.

https://doi.org/10.1017/S026988891100021X Published online by Cambridge University Press

We also had several *ad hoc* meetings where we demonstrated the functionality of the DSW and its browsing and searching facilities were evaluated. We mostly got positive feedback, however, there were some comments about the usability of the SSE and browsing through ontologies. People find it difficult to browse or search the OMR without being familiar with the underlying ontology structure. However, this was expected since the DSW addresses expert users with knowledge of the represented domain.

8 Synopsis and prospects

Our work intends to define a framework for capturing invaluable expert knowledge, sometimes implicitly contained in 3D resources, that is mostly undocumented, and therefore hard to be reused or automated, with the goal of fostering the next generation of semantically enabled systems and services. This work is based on an ontology-based management framework for digital shapes, which aims to address the need of a new approach to retrieve shapes, tools and publications related to 3D resources. We believe that the addition of explicit semantics to 3D content can significantly improve search results.

In this paper the OMR has been presented in detail that constitutes a first step in creating a large-scale framework for promoting the use of knowledge technologies for managing, searching and retrieving 3D media content. The primary goal of the OMR, which is part of the DSW, is the formalization and sharing of knowledge about 3D resources and their applications. The need for such a framework seems to grow exponentially as was discussed in the introductory and motivation sections of this paper. It is also evident that this need requires new approaches for retrieving information on shapes, but it also calls for an altogether different representation of shapes; one that takes into account not only the geometrical or structural properties of these objects, but also deals with their semantics as well.

None of the existing multimedia ontologies were sufficient enough to cover the representation of our 3D model needs (as described in Section 4), therefore we developed our own set of ontologies (common and domain-specific ontologies). We decided to use the OWL for modeling the knowledge in 3D objects because it is the most commonly used standard for ontological representation. There are currently plenty of tools and APIs supporting ontology development in OWL and can be easily coupled with many existing DL reasoners. Work is being done toward this direction with other ontology representation languages. For example, there are proposals of how to encode the CIDOC-CRM in OWL DL.

Effective 3D search and retrieval is a major challenge. For instance, a next generation search mechanism for 3D content could integrate content-based, geometry-driven criteria with concept-based, semantic-driven ones (see Figure 8). This search engine can be based on an innovative query formulation that supports geometric, semantic and combined search modalities. For example, it could be possible to pose queries such as find the 3D models in the repository that represent a vase with handles, and whose handles are globally similar in shape to a given query model. In the example, vase and handle could refer to semantic annotations and be resolved via a semantic search, whereas handles are globally similar in shape could be resolved by applying a geometric search to the models selected by the semantic search. This kind of query formulation interface would provide the unprecedented capability of capturing the intuitive notion of similarity expressed by the user in the query session.

Although text and image digitization and management are rather mature technologies, the technology necessary to benefit from 3D knowledge management is not yet fully exploited. There are still many areas in which further progress in research and consolidation of approaches is required. Methods and tools are required for making digital shapes machine-understandable and not just human-understandable as they are today, by developing semantic mark-up of content, in addition to intelligent agents and ontology infrastructures for 3D content (Bustos *et al.*, 2007).

Semantically enriched descriptions, indexing, and retrieval will allow the deep integration of 3D content into Digital Libraries, and fascinating new applications can be envisioned. Intelligent 3D

https://doi.org/10.1017/S026988891100021X Published online by Cambridge University Press

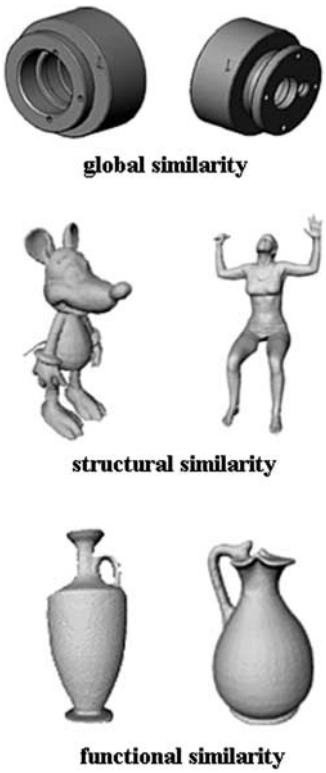


Figure 8 [Courtesy of IMATI-CNR] Retrieval of 3D content by combining geometric, structural and semantic criteria

acquisition will automatically segment and interpret any scanned objects, semantically enriching the data. New possibilities to work with 3D data will emerge, recognizing not only the low-level shape features, but also the structure and semantics. Once 3D editors are made aware of semantic properties of the models, the composition of new 3D objects based on existing content will be greatly simplified. 3D search engines will become highly intelligent, utilizing the available semantic shape analysis methods and the many different similarity notions on all conceptual levels.

Current limitations are mostly related to methodologies and the lack of specialized tools for 3D knowledge management. The areas in which progress would be particularly important and might be achieved in the relatively short term are:

(a) Developing solutions for improved reuse of 3D datasets by end-users. There is a need to support better an easy reuse of 3D objects as well as their use in a variety of contexts.
(b) Developing easy-to-use authoring tools for 3D knowledge management. Digitizing objects in 3D format is only a first step. The creation of complex 3D environments for end-user interaction will require the involvement of knowledge-management technologies. However, it would be beneficial to provide 3D professionals with an easy-to-use toolkit.
(c) Moving toward non-textual, semantic documentation and processing of 3D content. With the growth in 3D content there is a pressing need for progress in the field of semantic processing of such content. Development of ontologies is crucial to 3D content categorization, indexing, detailed search and retrieval and enhanced access.
(d) Undertaking ontological modeling to establish a broader base of domain knowledge, which is systematically encoded.
(e) Developing tools for the integrated management of the metadata related to 3D objects. There are specific aspects of this metadata that are unique to the digital domain including the digital object structure and parameters, the relationship of the digital object to the physical object

https://doi.org/10.1017/S026988891100021X Published online by Cambridge University Press

(e.g. coordinate systems and metrics); the digital provenance (i.e. processes used to create the model from the original data, tools and parameters employed, etc.). The digital provenance is a vital part of the long-term usability of the digital object, providing information for migration to new formats in the future. These tools would be based on the ontological framework mentioned earlier.

(f) Developing solutions for improved search. There is a need to better support the end-user in identifying and re-using 3D digital objects. Search and retrieval will be used in a variety of contexts with different search criteria.

Future directions of the work presented here will be focused on the following challenges:

1. provide Web Services for shape processing (e.g. smoothing, simplification, enhancement, etc.) and capture the knowledge related to the history of a digital shape (i.e. workflow information, software tools or algorithms used, parameterization and input information, etc.);
2. facilitate automatic, semantic annotation of 3D models;
3. enhance data repositories to exploit and reuse semantic annotations;
4. support intelligent discovery and retrieval of 3D models by improving the efficiency of semantic search engines.

Acknowledgments

This work was carried out within the scope of the AIM@SHAPE Network of Excellence supported by the European Commission Contract IST no. 506766. The authors would like to thank all the AIM@SHAPE partners. The authors were partially supported by the FOCUS K3D Coordination Action, EC Contract ICT-2007.4.2 no. 214993.

References

3D-COFORM. 2008. *Tools and expertise for 3d collection formation, FP7-IST Collaborative project*. http://www.3d-coform.eu

AceMedia. 2005. *Towards a common multimedia ontology framework report, FP6-IST Integrated Project*. http://www.acemedia.org

Aduna Open Source project. 2009. *Sesame: an open source RDF framework*. http://www.openrdf.org

AIM@SHAPE. 2005. *Advanced and innovative models and tools for the development of semantic-based systems for handling, acquiring, and processing knowledge embedded in multidimensional digital objects, FP6-IST Network of Excellence*. http://www.aimatshape.net

AIM@SHAPE. 2006a. *The Digital Shape Workbench*. http://dsw.aimatshape.net

AIM@SHAPE. 2006b. *Shape Repository*. http://shapes.aimatshape.net

Albertoni, A., Papaleo, L., Robbiano, F. & Spagnuolo, M. 2006. Towards a conceptualization for shape acquisition and processing. In *Proceedings of the 1st Workshop on Shape and Semantic*, Matsushima, Japan, 85–90.

Arnaud, R. & Barnes, M. 2006. *Collada: Sailing the Gulf of 3D Digital Content Creation*. AK Peters Ltd.

Arndt, R., Troncy, R., Staab, S., Hardman, L. & Vacura, M. 2007. Comm: designing a well founded multimedia ontology for the web. In *Proceedings of the 6th International and 2nd Asian Semantic Web Conference (ISWC2007+ASWC2007)*, Busan, Korea, November 11–15, 2007. 29–42. http://data.semanticweb.org/conference/iswc-aswc/2007/tracks/research/papers/29

ATK project. 2008. *3store: an RDF triplestore*. http://threestore.sourceforge.net

Baader, F., Calvanese, D., McGuinness, D., Nardi, D. & Patel-Schneider, P. (eds). 2003.*The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press.

Babik, M. & Hluchy, L. 2008. A testing framework for OWL-DL reasoning. In *Proceedings of the 4th International Conference on Semantics, Knowledge and Grid*, Beijing, China, December 3–5, 2008. 42–48.

Bloehdorn, S., Simou, N., Tzouvaras, V., Petridis, K., Handschuh, S., Avrithis, Y., Kompatsiaris, I., Staab, S. & Strintzis, M. G. 2004. Knowledge representation for semantic multimedia content analysis and reasoning. In *Proceedings of the European Workshop on the Integration of Knowledge, Semantics and Digital Media Technology (EWIMT)*, London, UK, November 25–26, 2004. http://www.image.ece.ntua.gr/publications.php

https://doi.org/10.1017/S026988891100021X Published online by Cambridge University Press

Bustos, B., Fellner, D., Havemann, S., Keim, D., Saupe, D. & Schreck, T. 2007. Foundations of 3D digital libraries—current approaches and urgent research challenges. In *Proceedings of the 1st International Workshop on Digital Libraries Foundations DLF07)*, Vancouver, Canada, June 23, 2007. 7–12. http://www.delos.info/index.php?option=com_content&task=view&id=576&Itemid=313

Catalano, C., Camossi, E., Ferrandes, R., Cheutet, V. & Sevilmis, N. 2008. A product design ontology for enhancing shape processing in design workflows. *Journal of Intelligent Manufacturing* **20**(5), 553–567.

Catton, C., Sparks, S. & Shotton, D. 2005. The imagestore ontology and the bioimage database: Semantic Web tools for biological research images. In *Proceedings of The International Conference on Computer as a Tool (EUROCON 2005)*, Belgrade, Serbia & Montenegro, November 21–24, 2005. 257–264.

Chen, D.-Y., Tian, X.-P., Shen, Y.-T. & Ouhyoung, M. 2003. On visual similarity based 3D model retrieval. *Computer Graphics Forum* **22**(3), 223–232.

COLLADA. 2009. *3D Asset exchange schema*. http://www.khronos.org/collada

Daras, P., Zarpalas, D., Tzovaras, D. & Strintzis, M. G. 2004. Shape matching using the 3D radon transform. In *Proceedings of the International Symposium on 3D Data Processing Visualization and Transmission*, (3DPVT 2004), Thessaloniki, Greece, September 6–9, 2004. 953–960.

Doulaverakis, C., Kompatsiaris, Y. & Strintzis, M. 2005. Ontology-based access to multimedia cultural heritage collections the REACH project. In *Proceedings of The International Conference on Computer as a Tool (EUROCON 2005)*, Belgrade, Serbia & Montenegro, November 21–24, 2005. 151–154.

Dublin Core 2006. *Dublin core metadata initiative*. http://dublincore.org

EPOCH. 2008. *Excellence in processing open cultural heritage, FP6-IST Network of Excellence*. http://www.epoch-net.org

Falcidieno, B., Spagnuolo, M., Alliez, P., Quak, E., Houstis, C. & Vavalis, M. 2004. Towards the semantics of digital shapes: the AIM@SHAPE approach. In *Proceedings of the European Workshop on the Integration of Knowledge, Semantic and Digital Media Technologies EWIMT 2004*, November 25–26, 2004. IEE.

Fikes, R., es2, P. & Horrocks, I. 2003. *OWL-QL—a language for deductive query answering on the Semantic Web*. Technical Report KSL-03-14, Knowledge Systems Laboratory, Computer Science Department, Stanford University. ftp://ftp.cs.pitt.edu/chang/revs/c916.pdf

FOCUS K3D. 2008. *FOCUS K3D, FP7-IST coordination action*. http://www.focusk3d.eu

Funkhouser, T., Kazhdan, M., Min, P. & Shilane, P. 2005. Shape-based retrieval and analysis of 3d models. *Communications of the ACM* **48**, 58–64.

Galinski, J., Kaya, A. & Moller, R. 2005. Development of a server to support the formal Semantic Web query language OWL-QL. In *Proceedings of the International Workshop on Description Logics*, Horrocks, I., Sattler, U. & Wolter, F. (eds). Edinburgh, UK, July 26–28, 2005. Available online as CEUR Workshop Proceedings Volume 147. http://ceur-ws.org/Vol-147/

Gutirrez, M., Garca-Rojas, A., Thalmann, D., Vexo, F., Moccozet, L., Magnenat-Thalmann, N., Mortara, M. & Spagnuolo, M. 2007. An ontology of virtual humans: incorporating semantics into human shapes. *The Visual Computer* **23**, 207–218.

Hendler, J. 2003. Science and the Semantic Web. *Science* **299**, 520–521. http://dx.doi.org/10.1126/science.1078874

Humanoid animation. 2009. *H-anim*. http://h-anim.org/Specifications/H-Anim200x

Hunter, J. 2001. Adding multimedia to the Semantic Web—building an mpeg-7 ontology. In *Proceedings of the International Semantic Web Working Symposium (SWWS)*, Stanford, USA, July 30–August 1, 2001. 261–281.

Hunter, J., Koopman, B. & Sledge, J. 2003. Software tools for indigenous knowledge management. In *Proceedings of the Museums and the Web*, Charlotte, 19–22 March 2003. http://www.archimuse.com/mw2003/papers/hunter/hunter.html

IBM Corporation. 2004. *IODT: integrated ontology development toolkit*. http://www.alphaworks.ibm.com/tech/semanticstk

International Council of Museums. 2009. *CIDOC Conceptual Reference Model*. http://cidoc.ics.forth.gr/

ISO FDIS 21127. 2007. *Information and documentation—a reference ontology for the interchange of cultural heritage information*. http://www.iso.org/iso/catalogue_detail.htm?csnumber=34424

ISO-IEC Moving Picture Experts Group working group. 2004. *MPEG-7 overview*. http://www.chiariglione.org/mpeg/standards/mpeg-7/mpeg-7.htm

Jiantao, P. & Ramani, K. 2007. An integrated 2D and 3D shape-based search framework and applications. *Computer-Aided Design and Applications* **4**(6), 817–826.

Kazhdan, M., Funkhouser, T. & Min, P. 2004. A comparison of text and shape matching for retrieval of online 3d models. Research and Advanced Technology for Digital Libraries. Lecture Notes in Computer Science. *Springer* **3232**, 209–220.

https://doi.org/10.1017/S026988891100021X Published online by Cambridge University Press

Kiryakov, A., Ognyanov, D. & Manov, D. 2005. OWLIM a pragmatic semantic repository for OWL. In *Proceedings of the International Workshop on Scalable Semantic Web Knowledge Base Systems (SSWS 2005)*, Lecture Notes in Computer Science. Springer-Verlag, **3807**, 182–192.

Lagoze, C. & Hunter, J. 2001. The abc ontology and model. In *Proceedings of the International Conference on Dublin Core and Metadata Applications 2001*, 160–176. National Institute of Informatics, Tokyo, Japan.

Lee, R. 2004. *Scalability report on triple store applications*. http://simile.mit.edu/reports/stores

Leifman, G., Meir, R. & Tal, A. 2005. Semantic-oriented 3d shape retrieval using relevance feedback. *The Visual Computer* **21**, 865–875.

McBride, B. 2002. Jena: a Semantic Web toolkit. *IEEE Internet Computing* **6**(6), 55–59.

Min, P. 2004. *A 3D Model Search Engine*. PhD thesis, Princeton University.

Noy, N., Fergerson, R. & Musen, M. 2000. The knowledge model of Protege-2000: Combining interoperability and flexibility. In *Proceedings of the 12th European Workshop on Knowledge Acquisition, Modeling and Management, EKAW 00*, 17–32. Springer-Verlag.

Pereira, F. 2001. The MPEG-21 Standard: Why an open multimedia framework? Interactive Distributed Multimedia Systems. Lecture Notes in Computer Science. *Springer* **2158**, 219–220.

Racer Systems GmbH & Co. (n.d.). *Renamed abox and concept expression reasoner RACER*. http://www.racer-systems.com

RDFS. 2009. *Resource description framework schema*. http://www.w3.org/TR/rdf-schema

SALERO. 2006. *Semantic audiovisual entertainment reusable objects, FP6-IST STREP*. http://www.salero.info

SCHEMA. 2002. *Content-based semantic sense analysis and information retrieval, FP6-IST Network of Excellence*. http://www.iti.gr/SCHEMA

SCULPTEUR. 2005. *Semantic and content-based multimedia exploitation for European benefit, FP7-IST STREP*. http://www.sculpteurweb.org

Seaborne, A. 2004. *RDQL—a query language for RDF*. http://www.w3.org/Submission/2004/SUBM-RDQL-20040109

Steinberg, D., Budinsky, F., Paternostro, M. & Merks, E. 2008. *EMF: Eclipse Modeling Framework*, 2nd edn. Addison-Wesley Professional.

Vasilakis, G., Pitikakis, M., Vavalis, M. & Houstis, C. 2005. A semantic based search engine for 3D shapes: design and early prototype implementation. In *Proceedings of the 2nd European Workshop on the Integration of Knowledge, Semantics and Digital Media Technologies*, London, UK, 391–397.

Vasilakis, G., Garcia-Rojas, A., Papaleo, L., Catalano, C. E., Robbiano, F., Spagnuolo, M., Vavalis, M. & Pitikakis, M. 2010. Knowledge-based representation of 3D shapes. *International Journal of Software Engineering and Knowledge Engineering* **20**(5), 739–760.

Veltkamp, R. C. 2001. Shape matching: similarity measures and algorithms. In *Proceedings of Shape Modelling International (SMI 2001)*, Genoa, Italy, May 7–11, 2001. 188–197.

VICTORY. 2006. *VICTORY, FP6-IST STREP*. http://www.victory-eu.org

W3C RDF Data Access Working Group. 2009. *Joseki: a SPARQL server for Jena*. http://www.joseki.org

W3C Recommendation. 2004. *OWL web ontology language*. http://www.w3.org/TR/owl-features

W3C Recommendation. 2008. *SPARQL Query Language for RDF*. http://www.w3.org/TR/rdf-sparql-query

X3D. 2009. *Extensible 3D*. http://www.web3d.org/x3d

https://doi.org/10.1017/S026988891100021X Published online by Cambridge University Press