Interoperable Multimedia Metadata through Similarity-based Semantic Web Service Discovery

Stefan Dietze[1], Neil Benn[1], John Domingue[1], Alex Conconi[2], Fabio Cattaneo[2],

[1] Knowledge Media Institute, The Open University, MK7 6AA, Milton Keynes, UK
{s.dietze, n.j.l.benn, j.b.domingue}@open.ac.uk
[2] TXT eSolutions, Via Frigia 27, 20126 Milano, Italy
{alex.conconi, fabio.cattaneo}@txt.it

Abstract. The increasing availability of multimedia (MM) resources, Web services as well as content, on the Web raises the need to automatically discover and process resources out of distributed repositories. However, the heterogeneity of applied metadata schemas and vocabularies – ranging from XML-based schemas such as MPEG-7 to formal knowledge representation approaches – raises interoperability problems. To enable MM metadata interoperability by means of automated similarity-computation, we propose a hybrid representation approach which combines symbolic MM metadata representations with a grounding in so-called Conceptual Spaces (CS). In that, we enable automatic computation of similarities across distinct metadata vocabularies and schemas in terms of spatial distances in shared CS. Moreover, such a vector-based approach is particularly well suited to represent MM metadata, given that a majority of MM parameters is provided in terms of quantified metrics. To prove the feasibility of our approach, we provide a prototypical implementation facilitating similarity-based discovery of publicly available MM services, aiming at federated MM content retrieval out of heterogeneous repositories.

Keywords: Semantic Web Services, Multimedia, Metadata, Vector Spaces.

1 Introduction

A continuously increasing amount of *digital multimedia (MM) content* is available on the Web, ranging from user-generated video content, commercial Video on Demand (VoD) portfolios to a broad range of streaming and IPTV resources and corresponding metadata records [19]. Besides, it became common practice throughout the last decade, to expose all sorts of MM content and metadata stored in one particular repository through a set of *Web services*, which provide Web-based access to software functionalities processing MM content and metadata, i.e. to retrieve, transcode or scale MM assets [21]. In line with the increasing usage of the term Web service in a broader sense, in the following we will use it synonymous with any kind of software functionality which is accessible through HTTP or any other IP-based layer, ranging from rather light-weight APIs, REST-based interfaces or standard Web service technology such as SOAP [22], UDDI [23] and WSDL [24].

Hence, the increasing accessibility of distributed *MM resources* – content as well as services – raises the need to automatically discover and compose

distributed content. In that, the highly heterogeneous nature of MM resources distributed across distinct repositories leads to the following key challenges:

C1.   Discovery of distributed MM services.
C2.   Discovery of distributed MM content.

However, w.r.t. these goals, several issues apply:

Concurrent metadata schemes and vocabularies. Distinct approaches to metadata representation do exist, ranging from light-weight tagging approaches as deployed within user-driven websites such as youtube[1] and general-purpose metadata standards such as Dublin Core [5] to fully-fledged domain-specific metadata standards such as MPEG-7 [10]. Besides, concurrent vocabularies – differing in terminology, syntax or language - are widely used to provide metadata records leading to further heterogeneities and ambiguities [11][19]. This issue also applies to Web service metadata provided based on syntactic descriptions such as WSDL [24] or semantic annotations based on OWL-S [12] or WSMO [25].

Lack of metadata comprehensibility and semantic meaningfulness. Metadata records lack expressivity due to merely syntactic annotations – usually based on XML schemas – not exploiting semantics of used structures and terminologies [1] [9][20]. In addition, current MM metadata schemas usually focus on the low-level parameters describing the actual format and audio-visual characteristics of MM assets, although a combined representation of both the actual content as well as its audio-visual format is required [14]. Moreover, even approaches such as [18] which exploit formal semantic representations, e.g. based on Semantic Web (SW) technologies such as OWL[2] or RDF-S[3], rely on either the common agreement on a shared conceptualisation or the formal representation of mappings, what is costly and error-prone. These issues hinder the automatic composition and processing of MM metadata and resources, and hence, do lead to interoperability issues.

Lack of rather fuzzy matchmaking approaches. Current approaches to match between a certain request and available MM resources usually perform strict one-to-one matchmaking and require the subscription to a certain vocabulary from both providers and consumers. In that, only resources from a highly limited number of repositories which represent an exact match with the requested parameters are retrieved, while similar and otherwise related resources which potentially are useful are being left aside.

Consequently, in order to enable interoperability between heterogeneous MM resource metadata, representation approaches are required which are meaningful enough to implicitly infer about inherent similarities across concurrent sets of MM annotations. In previous work [4], the authors proposed a representational approach combining symbolic knowledge representation mechanisms – as used by current MM resource metadata approaches – and SW technologies, with a representation in so-called *Conceptual Spaces (CS)* [8]. The latter consider the representation of knowledge entities, such as the ones described in MM metadata, through

---

[1] http://www.youtube.com

[2] http://www.w3.org/OWL/

[3] http://www.w3.org/RDFS/

geometrical vector spaces where measurable quality criteria represent individual dimensions. Particular metadata records, i.e. instances, are represented as members, i.e. particular vectors, in a CS what facilitates computation of similarities by means of spatial distances.

Here, we propose the application of our hybrid representational approach to model metadata of MM resources – i.e. MM services and MM content – in order to enable computation of similarities across heterogeneous repositories. In particular, low-level audio-visual MM characteristics, which are usually described by means of quantified attributes based on certain metrics, such as the MPEG-7 [10] descriptors *Dominant Color* or *Homogeneous Texture*, lend themselves to being represented in terms of vectors. Consequently, our hybrid representational approach appears to be well suited and hence, qualifies well to tackle MM metadata interoperability.

The remaining paper is organized as follows. We provide an overview on related work in the area of MM service and content metadata intereoperability in Section 2. Our approach to represent MM metadata is introduced in Section 3 followed by a prototypical application utilising our approach for similarity-based MM resource discovery in Section 4. Section 5 concludes and discusses our work.

2   Related Work

To satisfy the content need of a specific consumer, a federated MM content provisioning engine needs to discover (*C1*), the appropriate MM services (i.e. repositories) and (*C2*), the appropriate content. The following figure depicts this vision:
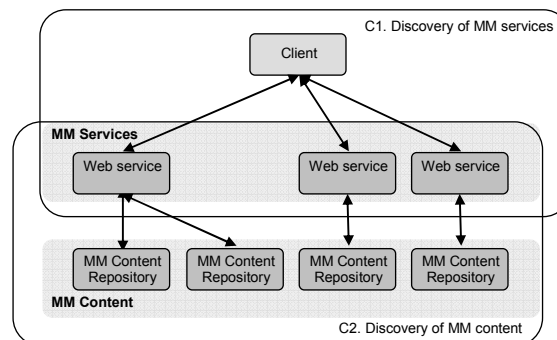
Fig. 1. Discovery of distributed MM services and content.

Given that both MM services and content utilize particular metadata vocabularies and schemas, approaching *C1* and *C2* requires taking into account related works from the areas of MM service metadata as well as MM content metadata interoperability.

2.1. MM Service Discovery through Semantic Web Services

With respect to *C1*, Semantic Web Services (SWS) technology aims at the automatic discovery, orchestration and invocation of distributed services on the basis of comprehensive semantic descriptions. SWS are supported through representation standards such as WSMO [25] and OWL-S [12]. We particularly refer to the *Web Service Modelling Ontology* (WSMO), an established SWS reference ontology and framework. WSMO is currently supported through dedicated reasoners, such as the Internet Reasoning Service IRS-III [2] and WSMX [26], which act as broker environments for SWS. In that, a SWS broker mediates between a service requester and one or more service providers. Based on a client request, the reasoner disovers potentially relevant SWS, invokes selected services and mediates potential mismatches.

However, the domain-independent nature of SWS reference models requires their derivation to facilitate the representation of certain domain-specific contexts. While SWS aim at automatic discovery of distributed Web services based on semantic metadata, current approaches usually rely on either the subscription to a common vocabulary and schema – i.e. a common domain ontology – or the manual definition of mappings between distinct service ontologies. In that, the previously introduced issues (Section 1) also apply to SWS technologies, demanding for approaches to deal with heterogeneities between distributed SWS. In that, approaches such as [15] aim at addressing the interoperability issue partially by resolving heterogeneities based on mapping approaches. For instance, [27] provides an attempt to support similarity detection for mediation within SWS composition by exploiting syntactic similarities between SWS representations. However, it can be stated that current approaches rely on the definition of a priori mappings, the agreement of a shared vocabulary or the exploitation of semi-automatic ontology mapping approaches. Hence, providing a more generic solution to automatically resolve heterogeneities between heterogeneous SWS remains a central challenge.

2.2. MM Metadata Interoperability

With respect to *C2*, a broad variety of research aims at interoperability between distributed MM (content) metadata. In general, the need for enriching non-semantic MM metadata through formal semantics is widely accepted [16] to enable more comprehensive query and retrieval facilities. For instance, [19] proposes an approach to semantically enrich MPEG-7 and TV-Anytime metadata through formal semantic expressions. In addition, [18] provides a way of formally expressing semantics of MPEG-7 profiles. While increasing the expressiveness of MPEG-7 based metadata, this work is limited to MPEG-7 exclusively. This also applies to the work proposed in [7], which provides an OWL expression of the MPEG-7 information model. In [9], the author provides a core ontology to annotate MM content to address interoperability. However, this approach relies on the subscription to a common vocabulary/schema – i.e. the suggested core ontology – what is not feasible in Web-scale scenarios. An entirely MPEG-21 based approach

for interoperable MM communication is proposed in [17]. The need to automatically discover and compose Web services to enable processing of MM content is expressed in [21], where the authors propose an approach based on OWL-S. However, the interoperability issues between heterogeneous symbolic service annotations (Section 2.1) also apply here.

While several approaches try to tackle MM metadata interoperability, it can be stated that the current state of the art usually relies on subscription to common (upper-level) vocabularies/schemas or the manual definition of mappings. Hence, issues arise when attempting to apply such approaches in Web-scale scenarios [11]. Therefore, analogous to the field of MM service annotations (Section 2.1), we claim that methodologies are required which allow for a more flexible alignment of distinct vocabularies.

3 Approach

With respect to the previously introduced issues (Sections 1 and 2), we claim that basing MM metadata representations on merely symbolic representations does not fully enable semantic meaningfulness [4] and hence, limits automatic identification of similarities across distinct schemas and vocabularies. In order to enable interoperability between heterogeneous MM resource metadata – representing MM content or services – representation approaches are required which are semantically meaningful enough to implicitly infer about inherent similarities. In that, we argue that a refinement of symbolic MM metadata through so-called Conceptual Spaces (CS) is better suited to overcome interoperability issues. While previous work [3][4] has shown that this approach can be applied to support interoperability between ontologies, here we apply it to facilitate interoperability between MM services and repositories.

3.1. Grounding MM Metadata in multiple Conceptual Spaces

We propose a two-fold representational approach – combining MM domain ontologies with corresponding representations based on multiple CS – to enable (a) similarity computation across concurrent MM metadata schemas and vocabularies and (b) the conjoint representation of low-level audio-visual features and the content semantics.

In that, we consider the representation of a set of n schema entities (concepts) E of a set of MM metadata records (ontology) O through a set of n Conceptual Spaces CS. Note, that we particularly foresee the application of this approach to metadata of both MM services as well as content. Schema entities in the case of MM content are, for instance, the MPEG-7 descriptors such as *Scalable Color* or *Edge Histogram*. In the case of MM services, a schema entity could be for instance a WSMO ontology concept. MM metadata values (instances) are represented as members, i.e. vectors, in the respective CS.

While still benefiting from implicit similarity information within a CS, our hybrid approach allows maintaining the advantages of symbolic MM metadata

representations and comprehensive domain ontologies, i.e. the ability to represent arbitrary relations and axioms. In order to be able to refine and represent ontological concepts within a CS, we formalised the CS model into an ontology [4]. Hence, a CS can simply be instantiated in order to represent a particular MM metadata schema entity. Referring to [8], we formalise a CS as a vector space defined through quality dimensions d_i. Each dimension is associated with a certain metric scale, e.g. ratio, interval or ordinal scale. To reflect the impact of a specific quality dimension on the entire CS, we consider a prominence value p for each dimension [8].

A particular member M – representing a particular value of a schema entity – in the CS is described through a vector defined by valued dimensions v_i. Following this vision, for instance the MPEG-7 schema entity Dominant Color could be represented through a CS defined by means of RGB values, where each of the spectrum colors represents one particular dimension of the CS. A certain shade of blue would then be represented through a member M_1, i.e. a vector with $M_1=\{(124, 177, 236)\}$.

Alignment between symbolic MM metadata representations and their corresponding CS (members) is achieved by referring the respective symbolic representation to the corresponding CS ontology containing the respective CS and member instances. In that, ontological MM metadata representations would import the CS ontology, while XML-based metadata could utilise a XML serialization of the CS ontology which is utilized as a particular controlled vocabulary. Hence, content and service semantics which are represented through particular domain ontologies are refined through CS to enable similarity-computation between distinct metadata sets.

3.2. Similarity-based Discovery of MM Resources

We define the semantic similarity between two members of a CS as a function of the Euclidean distance between the points representing each of the members. Hence, with respect to [4], given a CS definition CS and two members V and U, defined by v v_0, v_1, ...,v_n and u_1, u_2,...,u_n within CS, the distance between V and U is calculated as a normalised function of their Euclidean distance. For further details, please refer to [3][4].

In order to facilitate automated similarity computation between distinct MM metadata vocabularies and schemas, we provided a Web service (WS_{sim}) capable of computing similarities between multiple members in multiple CS. This Web service enables to automatically identify similarities between multiple MM metadata records, and hence, to automatically select the most appropriate (i.e. the most similar) MM metadata record for a given request. In that, given a set of MM metadata records, for instance based on formal semantics or XML, and a set of corresponding CS representations which refine the MM metadata schema and its values by means of vectors, WS_{sim} is able to compute similarities and consequently, to map and mediate between concurrent metadata schemas and vocabularies.

This Web service is provided with the actual MM metadata request R and the x MM metadata records MM_i that are potentially relevant for R:

$R \cup \{MM_1, MM_2, ..., MM_x\}$. *R* is provided as a set of measurements, i.e. vectors *{v₁..vₙ}* representing a set of *m* Members *M(R)* in available CS, which describe the desired metadata values, e.g. values measuring a certain MPEG-7 descriptor or the certain criteria describing MM Web service capabilities, such as a specific Quality of Service (QoS). Also, each MM_i contains a set of concepts (schema entities) *C={c₁..cₘ}* and instances (entity values) *I={i₁..iₙ}*. For each M_i within *R* the corresponding CS representations *CS={CS₁..CSₘ}* are retrieved by WS_{sim} from the available CS ontology [4]. Similarly, for each MM_j members *M(MMᵢ)* − which refine the instances of MM_j and are represented in one of the conceptual spaces *CS₁..CSₘ* − are retrieved: $CS \cup M(R) \cup \{M(MM_1), M(MM_2), ..., M(MM_x)\}$.

Based on the above ontological descriptions, for each member *vₗ* within *M(R)*, the Euclidean distances to any member of all *M(MMⱼ)* which is represented in the same space *CSⱼ* as *vₗ* are computed. Consequently, a set of *x* sets of distances is computed as *Dist(MMᵢ)={Dist(R,MM₁), Dist(R,MM₂) .. Dist(R,MMₓ)}* where each *Dist(R,MMⱼ)* contains a set of distances *{dist₁..distₙ}* where any *distₖ* represents the distance between one particular member *vᵢ* of *R* and one member refining one instance of the capabilities of *MMⱼ*. Hence, the overall similarity between the request *R* and any available *MMⱼ* could be defined as being reciprocal to the mean value of the individual distances between all instances of their respective capability descriptions:

$$Sim(R, MM_j) = \left(\overline{Dist(R, MM_j)} \right)^{-1} = \left(\frac{\sum_{k=1}^{n} (dist_k)}{n} \right)^{-1}$$

Finally, a set of *x* similarity values − computed as described above − which each indicates the similarity between the request *R* and one of the *x* available MM records MM_j is computed by WS_{sim}:

$Output(WS_{sim}) = \{Sim(R, MM_1), Sim(R, MM_2), .., Sim(R, MM_x)\}$.

As a result, the most similar MM_j, i.e. the closest MM record, can be selected and invoked. In order to ensure a certain degree of overlap between the actual request and the selected MM record, we also defined a *threshold similarity value T* which determines the minimum similarity which is required.

4 Application – Similarity-based Selection of Video Services

We provided a prototypical implementation which aims at similarity-based retrieval of public MM content. Note, that instead of applying the representational approach to individual MM content metadata, our prototypical application utilizes our approach to annotate MM (Web) services which operate on top of distributed MM content repositories. The available services were annotated following the representational approach proposed in Section 3.1. Hence, our proof-of-concept application facilitates similarity-based selection of MM services (i.e. *C1* in Section 1), which in turn process and retrieve MM content (*C2*). In that, federated retrieval

and processing of MM metadata is supported to facilitate interoperability. Our application makes use of standard SWS technology based on WSMO and IRS-III (Section 2.1) to achieve this vision.

The application dynamically discovers services which had been created in the context of the EC-funded project NoTube[4] and make use of the Youtube-API[5] as well as data feeds provided by BBC- Backstage[6] and Open Video[7].

4.1. Representing MM services through multiple CS

In fact, five different Web services had been provided, each able to retrieve content from distinct repositories through keyword-based searches. WS_1 is able to retrieve content from the Youtube channel of The Open University[8], while WS_2 provides Youtube content associated with the *entertainment* category following the Youtube vocabulary. WS_3 performs keyword-based searches on top of the Open Video repository, while WS_4 operates on top of the news metadata feeds provided by BBC Backstage. In addition, WS_5 provides Youtube content suitable for mobiles.
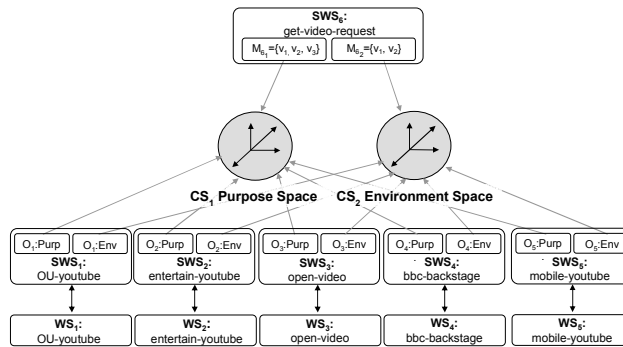

Fig. 3. MM service metadata refined in two distinct CS.

Based on the SWS reference model WSMO, we provided service annotations following the approach described in Section 3. In particular, we annotated the Web services in terms of the *purpose* they serve MM content for and the technical *environment* supported by the delivered content. In that, a simplified space (CS_1: *Purpose Space* in Figure 3) was defined to refine the notion of purpose by using three dimensions: *{((p₁∗information), (p₂∗education), (p₃∗leisure))}* = CS_1. The dimensions of CS_1 are measured on a ratio scale ranging from 0 to 100. For instance, a member P_1 in CS_1 described by vector *{(0, 100, 0)}* would indicate a rather educational purpose. In addition, a second space (CS_1: *Environment Space* in Figure 3) was provided to represent technical environments in terms of dimensions measuring the available resolution and bandwidth *{((p₁∗resolution),*

[4] http://projects.kmi.open.ac.uk/notube/

[5] http://code.google.com/intl/en/apis/youtube/

[6] http://backstage.bbc.co.uk/

[7] http://www.open-video.org/

[8] http://www.youtube.com/ou

$(p_2*bandwidth))$ = CS_2. For simplification, also the dimensions of CS_2 were ranked on a ratio scale. However, it is intended to refine the resolution dimension to apply an interval scale to both dimensions to be able to represent actual resolution and bandwidth measurements. Each dimension was ranked equally with a prominence of 1 in all cases.

By applying the representational approach proposed here, each concept of the involved heterogeneous SWS representations of the underlying MM services was refined as shared CS, while instances – used to define MM services and MM requests – were defined as members, i.e. vectors. In that, assumptions (*Ass*) of available MM services had been described independently in terms of simple conjunctions of instances which were individually refined as vectors in shared CS as shown in Table 1.

Table 1. Assumptions of involved SWS (requests) described as vectors in MS_1 and MS_2.

	Assumption $Ass_{SWSi} = (P_{1SWSi} \cup P_{2SWSi} \cup .. \cup P_{nSWSi}) \cup (E_{1SWSi} \cup E_{2SWSi} \cup .. \cup E_{mSWSi})$	
	Members P_i in CS_1 (purpose)	Members E_j in CS_2 (environment)
SWS_1	$P_{1(SWS1)}=\{(0, 100, 0)\}$	$E_{1(SWS1)}=\{(100, 100)\}$
SWS_2	$P_{1(SWS2)}=\{(0, 0, 100)\}$	$E_{1(SWS2)}=\{(100, 100)\}$
SWS_3	$P_{1(SWS3)}=\{(50, 50, 0)\}$	$E_{1(SWS3)}=\{(100, 100)\}$
SWS_4	$P_{1(SWS4)}=\{(100, 0, 0)\}$	$E_{1(SWS4)}=\{(100, 100)\}$
SWS_5	$P_{1(SWS5)}=\{(100, 0, 0)\}$ $P_{2(SWS5)}=\{(0, 100, 0)\}$	$E_{1(SWS5)}=\{(10, 10)\}$

Each service was associated with a set of members (vectors) in CS_1 and CS_2 to represent its purpose and the targeted environment. For instance, SWS_3 which provides resources from the Open Video repository, which in fact are of rather educational or information nature, was associated with a corresponding purpose vector *{(50, 50, 0)}*. While SWS_5 represents a Web service dedicated to MM content suitable for mobiles, a vector *{(10,10)}* indicating low resolution and bandwidth values was associated with SWS_5.
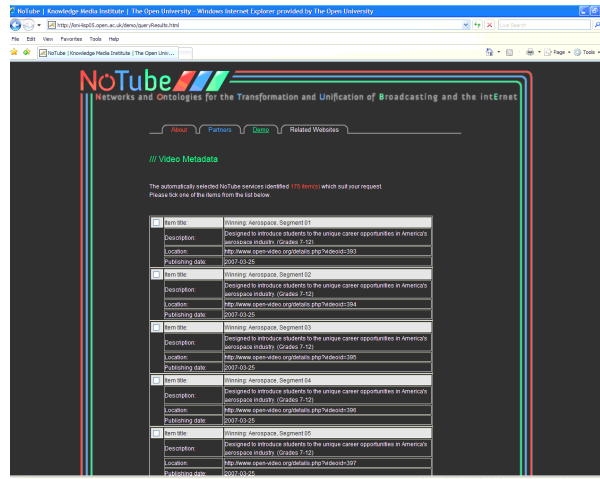

4.2. Similarity-based selection of MM services and content

An AJAX-based user interface (Fig, 4) was provided which allows users to define MM content requests by providing measurements describing their context, i.e. the purpose and environment, and search input parameters, i.e. a set of keywords. For instance, a user provides a request R with the search keyword "Aerospace" together with measurements which correspond to the following vectors: $P_1(R)=\{(60, 55, 5)\}$ in CS_1 and $P_2(R)=(95, 90)\}$ in CS_2. These vectors indicate the need for content which serves the need for education or information and which supports a rather high resolution environment.

Table 2. Automatically computed similarities between request R and available SWS.

	Similarities
SWS_1	0.023162405
SWS_2	0.014675636
SWS_3	0.08536871
SWS_4	0.02519804
SWS_5	0.01085659

Though no MM service matches these criteria exactly, at runtime similarities are calculated between R and the related SWS (SWS_1-SWS_5) through the similarity computation service WS_{sim} described in Section 3.2. This led to the calculation of the similarity values shown in Table 2. Given these similarities, our reasoning environment automatically selects the most similar MM service (SWS_3) and triggers its invocation.

Fig. 4. Screenshot of AJAX interface depicting MM metadata retrieved from the Open Video repository after similarity-based selection & invocation of MM services and metadata.

As illustrated above, our application utilises our representational mechanism (Section 3.1) to support similarity-based selection of distributed MM services. Hence, though just deploying our representational approach to MM services rather than MM content, our proof-of-concept prototype illustrates the applicability of our approach for similarity-based MM metadata discovery.

5 Conclusions

In order to facilitate interoperability between heterogeneous MM resources distributed across distinct repositories, we identified two major challenges – the discovery of appropriate MM services and the retrieval of the most appropriate MM content. However, addressing these challenges requires interoperability between concurrent metadata annotation schemas and vocabularies. To facilitate such interoperability, we proposed a two-fold representational approach. By representing MM annotation schema entities as dedicated vector spaces, i.e. CS, and corresponding values as vectors, similarities are computable by means of distance metrics. Our approach is realised through a dedicated CS ontology.

To prove the feasibility of our approach, we introduce a prototypical application which utilises our representational approach to support discovery of MM services across distributed MM repositories. As a result, we enable similarity-based

discovery of the most appropriate MM service for a given request, and hence, enable federated MM content and metadata searches across distributed repositories. However, while the current matchmaking algorithm considers instance similarity as exclusive suitability measure, future work will deal with the combined consideration of logical expressions and instance similarity.

We claim that our representational approach is particularly applicable to the domain of MM, where the majority of descriptors is based on quantified metrics, and hence, is well suited for metric-based representations such as vector spaces. The authors would like to highlight that providing the representations proposed here requires an additional effort, which needs to be investigated within future work. In this respect, please note that certain CS, for instance the one describing the notion of color, are reusable given that these are required for a variety of MM parameters. As another restriction, our approach foresees that distinct parties share common CS. However, given the wide-spread usage of upper-level ontologies such as DOLCE [6], SUMO [13] or OpenCyc[9] together with availability of common MM metadata standards [10] and ontologies [1][18], the agreement on common CS becomes increasingly applicable. Future work will be concerned with the evaluation of the effort required to utilise our representational model, and also, with carrying out further case studies.

References

[1] Arndt, R., Troncy, R., Staab, S., Hardman, L., Vacura, M., COMM: Designing a Well-Founded Multimedia Ontology for the Web. ISWC/ASWC 2007: 30-43.

[2] Cabral, L., Domingue, J., Galizia, S., Gugliotta, A., Norton, B., Tanasescu, V., Pedrinaci, C. (2006): IRS-III: A Broker for Semantic Web Services based Applications. In proceedings of the 5th International Semantic Web Conference (ISWC 2006), Athens, USA.

[3] Dietze, S., Gugliotta, A., Domingue, J., (2008) Conceptual Situation Spaces for Situation-Driven Processes. 5th European Semantic Web Conference (ESWC), Tenerife, Spain, June 2008.

[4] Dietze, S., and Domingue, J. (2008) Exploiting Conceptual Spaces for Ontology Integration, Workshop: Data Integration through Semantic Technology (DIST2008) Workshop at 3rd Asian Semantic Web Conference (ASWC) 2008, Bangkok, Thailand

[5] Dublin Core Metadata Initiative (2006). Dublin Core Metadata Terms. http://dublincore.org/documents/dcmi-terms/.

[6] Gangemi, A., Guarino, N., Masolo, C., Oltramari, A., Schneider, L. (2002). Sweetening Ontologies with DOLCE. In: A. Gómez-Pérez, V. Richard Benjamins (Eds.) Knowledge Engineering and Knowledge Management. Ontologies and the Semantic Web: 13th International Conference, EKAW 2002, Siguenza, Spain, October 1-4, 2002.

[7] Garcia, R., Celma, O. (2005) Semantic Integration and Retrieval of Multimedia Metadata, 2nd European Workshop on the Integration of Knowledge, Semantic and Digital Media (2005).

[8] Gärdenfors, P. (2000), Conceptual Spaces - The Geometry of Thought. MIT Press, 2000.

---

[9] http://www.opencyc.org/

[9] Hunter, J. Enhancing the semantic interoperability of multimedia through a core ontology. IEEE Trans. Circuits Syst. Video Techn. 13(1): 49-58 (2003)

[10] Moving Picture Experts Group, ISO/IEC JTC1/SC29/WG11 – MPEG-7, http://www.chiariglione.org/mpeg/standards/mpeg-7/mpeg-7.htm.

[11] Nack, F., van Ossenbruggen, J. & Hardman, L. (2005) That Obscure Object of Desire: Multimedia Metadata on the Web (Part II). IEEE MultiMedia, Vol 12, No. 1, pp.54 - 63.

[12] OWL-S 1.0. http://www.daml.org/services/owl-s/1.0/.

[13] Pease, A., Niles, I., Li, J.(2002), The suggested upper merged ontology: A large ontology for the semanticweb and its applications. In: AAAI-2002Workshop on Ontologies and the Semantic Web. Working Notes (2002).

[14] Petridis, K., Bloehdorn, S., Saathoff, C., Simou, N., Dasiopoulou, S., Tzouvaras, V., Handschuh, S., Avrithis, Y., Kompatsiaris, I., Staab, S. (2006). Knowledge Representation and Semantic Annotation of Multimedia Content, IEE Proceedings on Vision Image and Signal Processing, Special issue on Knowledge-Based Digital Media Processing, Vol. 153, No. 3, pp. 255-262, June 2006.

[15] Radetzki, U., Cremers, A. B. Iris: A framework for mediator-based composition of service-oriented software. In ICWS, pages 752–755. IEEE Computer Society, 2004.

[16] Stamou, G. van Ossenbruggen, J. Pan, J.Z. Schreiber, G. Smith, J.R. (2006), Multimedia annotations on the semantic Web, IEEE Multimedia, Volume: 13, Issue: 1, page(s): 86- 90, Jan.-March 2006

[17] Timmerer, C., Hellwagner, H. (2005). Interoperable Adaptive Multimedia Communication, IEEE Multimedia Magazine, vol. 12, no. 1, pp. 74-79, January-March 2005.

[18] Troncy, R., Bailer, W., Hausenblas, M., Hofmair, P., Schlatte, R.. Enabling Multimedia Metadata Interoperability by Defining Formal Semantics of MPEG-7 Profiles. In: 1st International Conference on Semantics And digital Media Technology (SAMT 06), Athens, Greece. December 2006.

[19] Tsinaraki C., Polydoros P., Kazasis F., Christodoulakis S., Ontology-based Semantic Indexing for MPEG-7 and TV-Anytime Audiovisual Content, Special issue of Multimedia Tools and Application Journal on Video Segmentation for Semantic Annotation and Transcoding, 2004.

[20] Van Ossenbruggen, J., Nack, F. & Hardman, L. (2004) That Obscure Object of Desire: Multimedia Metadata on the Web (Part I). IEEE MultiMedia, Vol 11, No. 4, pp.38-48.

[21] Wagner, M., Kellerer, W. 2004. Web services selection for distributed composition of multimedia content, Proceedings of the 12th annual ACM international conference on Multimedia, October 10-16, 2004, New York, NY, USA.

[22] W3C: Simple Object Access Protocol, SOAP, Version 1.2 Part 0: Primer, (2003). (http://www.w3.org/TR/soap12-part0/).

[23] W3C: Universal Description, Discovery and Integration: UDDI Spec Technical Committee Specification v. 3.0, (2003). http://uddi.org/pubs/uddi-v3.0.1-20031014.htm).

[24] W3C: WSDL: Web services Description Language (WSDL) 1.1, (2001). (http://www.w3.org/TR/2001/NOTE-wsdl-20010315

[25] WSMO Working Group (2004), D2v1.0: Web service Modeling Ontology (WSMO). Working Draft, (2004). (http://www.wsmo.org/2004/d2/v1.0/).

[26] WSMX Working Group (2007), The Web Service Modelling eXecution environment, http://www.wsmx.org/.

[27] Wu, Z., Ranabahu, A., Gomadam, K., Sheth, A.P., Miller, J.A., Automatic Composition of Semantic Web Services using Process Mediation, Proceedings of the 9th International Conference on Enterprise Information Systems (ICEIS'07), Funchal, Portugal Jun. 2007. pages 453-461.