Knowledge Representation in WERKL, an Architecture for Intelligent Multimedia Information Systems

Marcus Herzog[1] and **Paolo Petta**[2]

Abstract.
Multimedia information systems handle vast quantities of media resources. As a consequence, it is difficult to keep track of the semantic content of these items, especially if they were produced by different users of the system. We are interested in developing a formalism and corresponding tools that will be capable of abstracting concepts, ideas, and lines of thought expressed in the media by inferring relationships between the content of different resources. The tools will emphasize the role of the system as a partner augmenting the capabilities of the human user. As such they will tackle the problems of collaboration (human-to-human, human-to-machine), knowledge-sharing, and knowledge-retrieval in multimedia information systems. Users are not (indeed, cannot be) assumed to have complete knowledge of the dynamically changing content of the system; instead, they engage in an exploration of the available resources, during which they are offered opportunities to analyse their information need from different points of view.

1 INTRODUCTION

Computer-based information systems manage digital resources which represent properties of the covered domains. Typically, human decision-makers use these abstractions of the real world as foundations for their decision-making processes. Additionally, the systems may themselves apply algorithms exploiting the available domain knowledge to generate explicit representations of information that is hidden in the mass of data. In either case, information is generated by selecting a certain set of data in respect to a given information need.

Traditionally, these abstract resources were very limited in terms of expressiveness and user interaction capabilities. First approaches settled for file systems that were basically used solely to store alphanumeric data. Further research led to the development of database management systems (used primarily for structured data) and information retrieval systems (used for unstructured — mostly textual — data). These systems impose a certain interpretation and interaction paradigm on data in order to allow the formulation of information processing activities. Theories and models within these technologies are used as guidelines for developing information processing applications.

The rapid development of the information processor's workbench — the computer — in terms of speed and storage capabilities along with its connection to global networking gave rise to the desire to include more semantically expressive data in information systems to ease the communication of information. *Multimedia* is the term that nowadays denotes applications that handle pictures, drawings, animations, videos, audio, and other media types. *Multimedia information systems* use these media types to communicate information. They should offer functions to create, store, manage, retrieve, process, and use media objects. The change of the underpinning technologies entailed by the transition from traditional to multimedia information systems raises a vast amount of research issues, some of which we will discuss in more detail in the following sections.

In our view, one of the most important impacts of this new technology will be related to the novel opportunities for dealing with the problems encountered in *open-world domains*. In contrast to traditional "closed" information systems these emerging systems are "open" in respect to their problem domain which is typically very dynamic. Thus, it cannot be assumed that the structure of the encoded information is fixed and defineable a priori or in an exhaustive manner. Instead these systems have to be adaptive to changing situations including a more comprehensive active involvement of the human user into the information processing activity. This follows from the assumption that a large proportion of information needs typically occurring in open-world domains can be most adequately satisfied by an *explorative* rather than a *search* approach. In contrast to the search approach where the emphasis lies on finding exactly matching information objects, the explorative approach stresses the notion of investigating the space of possible interesting information objects via a guided dialog between user and system, thereby resulting in a "deeper" understanding of the problem domain.

In the following we will first shortly discuss WERKL, an architecture for intelligent multimedia information systems in open-world domains and its conceptual foundation, the decomposition of the information space into architectonic and semantic spaces. Building on this description of the basic architecture we will investigate knowledge representation issues connected to our specific architecture and to the field of multimedia information systems for open-world domains in general. We will also give example applications of the WERKL architecture and discuss related research before concluding the paper.

2 THE ARCHITECTURE OF WERKL

2.1 Architectonic and Semantic Spaces

The content of the information space of open-world domains can be qualified using the notion of architectonic and semantic spaces [24]: while the "indefinitely filled" semantic space covers the totality of all possible interpretations of the subject matter, architectonic spaces are explicit representations of particular views on the content, according to different information needs. The totality of all architectonic spaces

[1] CD-Lab for Expert Systems, Vienna University of Technology, A-1040 Vienna, Austria, Email: herzog@dbai.tuwien.ac.at
[2] Austrian Research Institute for Artificial Intelligence, A-1010 Vienna, Austria, Email: paolo@ai.univie.ac.at

ECAI 96 Workshop. *KRIMS 96*
Edited by G.Vouros, W.Klas, T.Sellis, M.Vazirgiannis

erected over a shared semantic space thus reflects all of the investigated interpretations and applied structuring principles, i.e. the uses made of the available information.

Open-world domains lack a clear boundary delimiting a relevant part of the information space. In practice this is reflected by a dynamically changing coverage: the semantic space is expanded as hitherto unconsidered subject areas become of relevance; similarly, the architectonic space is extended with the inclusion of novel points of view on the current semantic space. Systems aimed at providing user-support in such settings consequently have to take into consideration both sides of the coin: in order to provide effective assistance, an explicit and rich representation of viewpoints on the domain should be implemented. On the other hand, attention has to be paid not to constrain the required freedom of movement of users in the information space, and not to obstruct their perception of possible alternative interpretations. Such shifts of focus may occur between well-established perspectives, idiosyncratic positions reflecting personal preferences or original work, and finally the yet-to-be-formalized *semantic space* itself. As already mentioned, we use this term in the tradition of [54] and [11] for the concept of the "domain of possible expression . . . where meanings or interpretations come into existence" [24]. This notion of semantic space is contrasted to *architectonic spaces*, each of which represents one particular point of view and encodes the corresponding stable properties. This reification allows a shared manipulation of the content by both user and system.

A system providing adequate support for open-world domains thus should provide

- rich architectonic spaces and associated powerful methods for well-established points of view: these serve the twofold purpose of enabling complex and authoritative reasoning processes — e.g., traditional expert system functionalities — within these circumscribed areas, as well as providing points of reference for tasks extending beyond their scope;
- support for the construction of idiosyncratic architectonic spaces: methods for automatic indexing and classification of available data come with related advantages such as guaranteed completeness with respect to the available resources and the integration of any new sources of information that might become available. In addition to maintenance within a single architectonic space, relationships between different spaces can be exploited for tasks such as refinement, consistency checking, or construction, possibly based on inductive methods. This latter variant can also be used to notify users of possible opportunities, which brings up another class of functionalities, which provide support for the semi-automatic creation of new points of view: compared to the former methods, gains in flexibility are here obtained at the price of an increased maintenance work burden placed on the users;
- support for the investigation of the semantic space: given that the possibilities to provide an active direct support in this area are limited, an emphasis may be placed on alleviating users of management burdens and reducing encountered cognitive overheads — e.g., by including adequate task models and according user interface capabilities. In addition, there are various sources for possibly valuable clues and hints, such as supporting analyses of the interrelationships between existing architectonic spaces and of their grounding in semantic space, or using libraries of high-level specifications of possible courses of action.[3] However, at all times the abovementioned risks of an (inadvertent) introduction of biases or other hindrances to a broad access of the semantic space have to be considered.

[3] The growing corpus of such patterns is documented in conference series such as the annual conferences on Pattern Languages of Programs (PLoP).

It follows that the kind of assistance systems can provide varies with the momentary context: within architectonic spaces, substantial tasks can be delegated to the computer, which in addition can assume authoritative roles e.g., assessing the actions of the user. In contrast, during interactions involving the semantic space the cooperation becomes more of a collaborative kind, with the computer acting as critic that may be ignored or overruled. [52].

2.2 The WERKL Architecture

Following our understanding of the distinct problems found in representing open-world knowledge as specified by the dualism of semantic and architectonic spaces, the basic architecture of WERKL comprises two layers: a formal *index layer* and a semi-formal *data layer*. The formal index layer accounts for explicitly represented properties of the artifacts while the semi-formal data layer covers the totality of all potential associations tied to these elements.

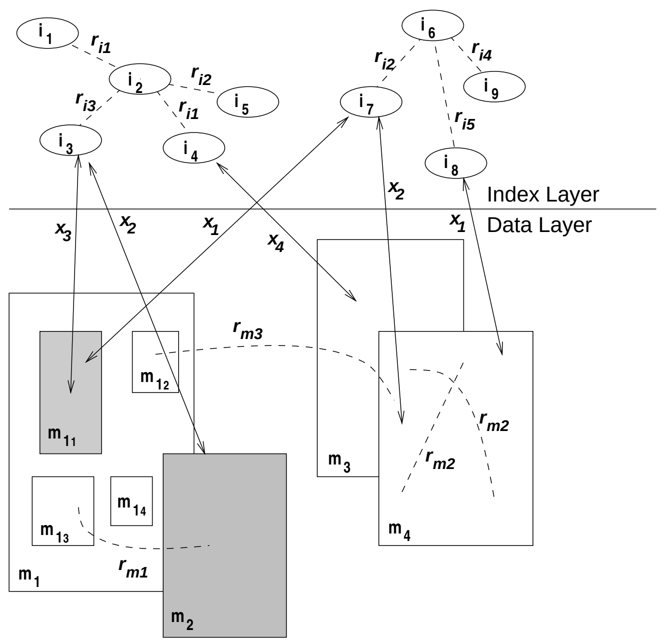

Figure 1. The WERKL system architecture

The system architecture (Figure 1) can be conceptually represented by a six-tuple $[M, R_m, I, R_i, X, O]$ where:

M is a set of media data;
R_m is a set of relations defined on M; } *data layer*
I is a set of index layer nodes that are used to qualify M;
R_i is a set of relations defined on members of I; } *index layer*
X is a set of relations defined on M and I;
O is a set of information generating computational operations defined on M, R_m, I, R_i, and X.

M is a collection of (multi)media data, e.g. portions of text, video chunks, audio clips, etc. The model is not limited to flat objects, but is also able to handle structured objects, i.e., aggregations of simple objects. The actual implementation of these objects can draw from the multitude of proposed standards to handle multimedia data.

R_m is a set of relations defined on objects of M. In the most simple case there is only one kind of relation, namely the association of two objects without any more detailed semantics. More advanced models can incorporate different types of relations such as temporal or spatial relationships, depending on the purpose of the application domain. The union of M and R_m resembles the notion of *hypermedia*.

Items in I attribute objects in M. The purpose of these items is to qualify the objects in M from a certain point of view, thus making semantic information explicit to the system and the user. Typically, I is structured by using relationships defined in R_i. R_i can hold domain-independent and/or domain-dependent relations. It is expected that different pairs $<I, R_i>$ will be developed in time even over the same set M, each of the I representing a distinguished point of view on the objects in M.

X is the non-empty set of relations defined between items of I and M. Computational operations in O exploit these relationships to facilitate intelligent behaviour of the system. We will discuss the different types of intelligent actions performed by the system and the corresponding knowledge structures and algorithms in the next section. For now we will focus on the *exploration* process as the central high-level mechanism within WERKL.

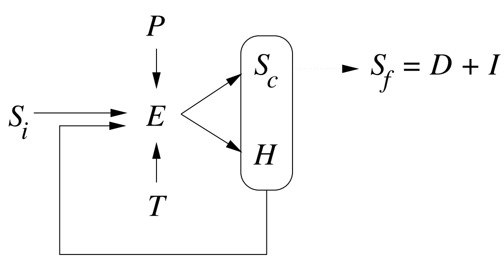

Figure 2. The WERKL exploration cycle

This process (Figure 2) can be represented by a six-tuple [P, T, $< S_i \ldots S_f >$, H, D, E],[4] where

- P is the union of M and I;
- T is a set of heuristics that compute the semantic distance between objects of M using R_i, R_m and X.
- $< S_i \ldots S_f >$ is a poset of property sets. These property sets are non-empty sets of instances of M and I. The sets are the intermediary results obtained during an exploration process, S_i being the initial set, S_f the final set that finally characterises the available information found that meets the users's information need, and the current one being referred to as S_c;
- H is the exploration history: it is a set of sequences of property sets and applied transformations;
- D is an initially empty set of data items meeting the requirements defined by S_c, $(D = S_c - I)$;
- E is one exploration cycle that extends D by taking P, T, S_c, and H as input and producing a series of transformations on S_c (leading to the next S_c and the corresponding next D).

[4] This formalization of the exploration process is an adaptation of the formalism introduced by Smithers and Troxell [48] for the design process.

We will now illustrate the formal constructs that we introduced using an example of a possible interaction scenario in one of our application fields, the HYSAT project, which is described in greater detail in a later section.

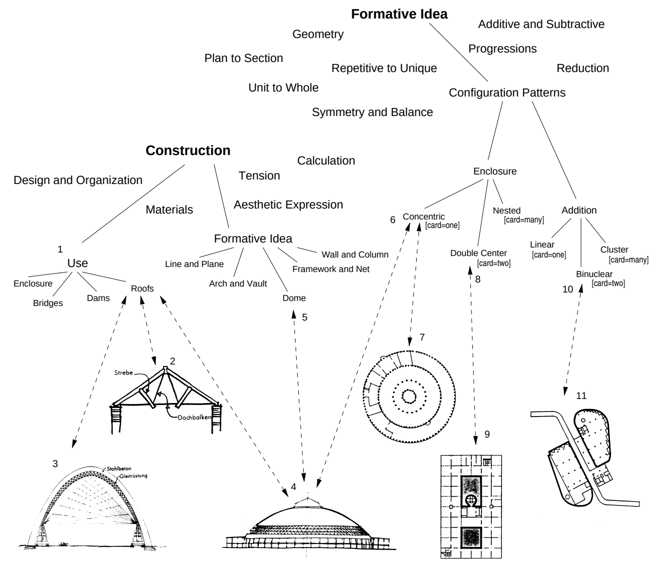

Figure 3. Exploration in HYSAT

Figure 3 visualises the conceptual entities within the HYSAT application: the term hierarchies depicted in the upper half of the graphic represent ontologies in the index layer and are instances of items of I and R_i. The graphic is of course just exemplary and should not be taken literally, as it shows only the items that are necessary for the further discussion. The drawings in the lower part symbolise media items in M. For reasons of simplicity we did not include items in R_m in the graphic. The dashed arrows between the terms and the drawings stand for relations in X, meaning the term attributes the content of the drawing.

As HYSAT is conceived as a design-supporting tool, the purpose of the information system is to locate documents of relevant design examples given a design problem. Design examples are used to communicate design knowledge. In the following we will look at a typical interaction cycle a designer will undertake working with the system. In this cycle, some steps are initiated by the system while others are governed by the user. The numbers (1 to 11) in Figure 3 state the chronological order of the interaction steps.

Let us assume the designer sets out with the question: "How to design a wide-spanned roof sheltering a gymnasium on a rectangular site?". She thus initiates the property set S_i with the term Roof. The system uses the set S_i and E to compute the set S_c, which then consists of design cases (multimedia items) referred to by the term Roof and all terms referring to these cases in the ontology Construction (1–3). By browsing the multimedia items she gets interested in an example of a dome-shaped roof (4). Because she wants to know the consequences of a dome-shaped roof for the ground plan, she seeks for a corresponding example in the ontology Formative Idea.[5] The descriptor in the ontology Construction for the selected roof is Dome (5). At this point, the user wants to change point of view

[5] Note the polysemy of the term "Formative Idea" in the example, i.e. its different use in different ontologies.

and asks the system to determine an alternative description exploiting operations in E[6] of the documents referred to by Dome. The proposed solution combines the concepts Roofs of the current ontology with the concept Concentric of the “Formative Idea” ontology. The user follows to this new perspective (6). Additional information items retrieved for Concentric (7) turn out to be inconsistent with the rectangular site prescribed in the problem definition.

To find alternative solutions, the user tries to weaken the requirement in the current point of view and the system discovers Double Center (8) as a closely related term concerning Enclosure. Examples prove to be suitable for a rectangular site (9), e.g. two domes. At this point, the user furthers the search by asking the system for alternatives to the found solution (which supersedes some of the interim assertions, such as the concept Dome). Using the *regularity*[7] [36] identified between the Enclosure and Addition hierarchies of the ontology Formative Idea, the system chooses the concept Binuclear (10) and proposes the documents referenced by it (11). These are finally accepted as solution to the problem. Note how the information retrieval process led to the exploration of regions of the design space — represented by selections of multimedia documents — that the user was previously unaware of.

Candidate techniques to implement the conceptual entities of the WERKL architecture will be discussed in the next section. The design of the system architecture is intended to allow for various approaches in regard to the actual implementation. This is necessary also because of the rapid development in research fields that focus on the supporting technologies. The intention is to provide a framework with this system architecture that can be used to realize different actual implementations in different domains.

3 KNOWLEDGE REPRESENTATION ISSUES

3.1 Overview

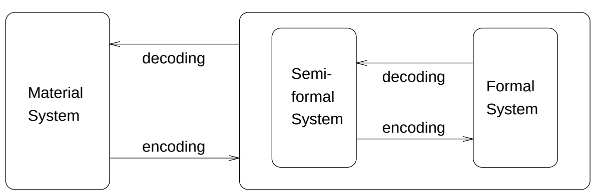

Figure 4. The two stage formalization process

The ultimate goal of information systems is to deal with information found in real world scenarios. The two-stage abstraction approach (see Figure 4) — via the data and index layer — alleviates the difficult task of capturing and formalizing real world information. This is especially true for domains where multimedia items can be used to supply a more immediate, perceptually accessible computer-based representation of such information: e.g., a picture can be represented more “comprehensively” by a scanned image than by some verbal description. In contrast to pure data acquisition methodologies this approach combines the simultaneous acquisition of both formal and informal aspects of the domain: e.g., the insertion of a picture along with a formal description of certain aspects of its content in an appropriate formalism. In terms of the example given above, while the picture holds “indefinitely rich” semantic information (i.e., the picture can be interpreted in an unlimited number of ways by humans), the formalized descriptions will always just cover a limited amount of this semantic content.

In the present case the incomplete coverage provided by the domain knowledge encoded in the index layer is arguably even more of an advantage than a drawback: the lack of a once-and-for-all commitment for structuring of the data warrants a flexible handling of the system’s content. If users feel that the system does not satisfy their needs in terms of expressiveness they can modify the index layer accordingly: the basic attitude is to provide as much freedom as reasonably possible. We see finding an appropriate balance as one of the crucial topics for future multimedia information systems.

In contrast to classical information systems that were primarily designed to perform retrieval of exactly matching items given a description of these items in some formal language, multimedia information systems tend to support a more interactive nature of the retrieval process. This is necessary as users are not assumed (indeed, cannot be) to have a comprehensive understanding of the content of the system. Instead of having to learn the structure of the index space, users should be to endorsed in an serendipitous exploration of the semantic space.

In order to meet all these goals, WERKL has to draw extensively on different kinds of knowledge during different stages of the information processing cycle:

- knowledge about the domain;
- knowledge about the task at hand;
- knowledge about the user.

During the acquisition process knowledge about the domain is used to guide the user in the construction process. Knowledge about the user’s task eases the human-computer interaction process during explorative use of the system as the system’s responses are tailored to the problem setting. Finally, user models will tune the performance of the system over time and personalize individual sessions. To allow for accumulation of knowledge, formal models are needed that function as generic classes of these knowledge types.

3.2 Techniques

Ontologies in a Federated Architecture

In a first approach, the index layer of WERKL, comprising domain knowledge and meta-knowledge about the data layer, will be encoded in ontologies, “explicit, partial accounts of a conceptualisation”, where a conceptualisation is “an intensional semantic structure which encodes the implicit rules constraining the structure of a piece of reality” [20, 21]. Regarding the domain knowledge, some ontologies will provide vocabularies (as e.g., in the Ontolingua ontologies for simple geometry or standard dimensions and units) while other ontologies will use these basic terms to encode the different interpretations of the contents of the data layer. Different ontologies are thus used to formalise different aspects of the semantic space by reifying abstract concepts implicitly represented in the underlying media (see e.g. [25] for a recent review of the kinds of metadata that have been used for different digital media types).

While the WERKL framework itself does not impose any restrictions on the kind of formalism to be used for the implementation of the index layer and the associated operations, a number of properties seem to make the use of description logics (DLs) a preferable choice over other knowledge representation formalisms. These properties concern both the local construction and maintenance of single

[6] E.g. by applying the Minimum Description Length principle [37].

[7] A structuring paradigm where the identification of mappings between hierarchies of different but related conceptual domains is used to support a number of inferences, also providing domain-oriented browsing strategies.

ontologies, as well as the support offered for interoperation in a distributed application scenario.

Viewed in a local context, there exists a body of well-understood theories along with methodologies and tools which can be readily exploited. The advantages following from the declarative approach are complemented with the availability of a thorough theoretical analysis of computational tractability and performance, including the possibility of a clear assessment of the benefits as well as the limitations of the approach. Opportunities in both maintenance of the knowledge body as well as the supported operations — including analogical reasoning; the automatic validation of integrity constraints; the discovery of new knowledge via induction; the organization and weakening of queries; the availability of incremental procedures such as query by refinement; semantic query optimization; dealing with incomplete or generic information; or the capability to provide intensional answers (e.g., [36, 27, 55]) — thus come together with an awareness of the limitations — regarding e.g. expressiveness[8] or possible tradeoffs between expressive power and inference properties [7], for which there also exist known remedies such as the inclusion of procedural "safety hatches" for the implementation of specifically required functionalities. Finally, there also exist criteria and procedures for the determination of the sufficiency of the expressive power of candidate DLs via the use of knowledge-specification languages (e.g. [19, 16]).

The catalogue of features speaking in favour of the deployment of DL is further extended in a significant way as the perspective is widened to span across multiple information sources, which in the context of WERKL is of relevance for the support of both individual viewpoints as well as change of perspectives between different interpretations of the semantic space. For individual viewpoints, being able to tap into other existing knowledge structures to identify new relevant information sources or promising candidates is an important feature. Analogously, being able to access remote viewpoints entails more than a merely quantitative extension of the support the system can provide. Virtually all of the existing efforts aimed at interoperability within information networks presuppose the use of a declarative representation of knowledge (or respectively the availability of facilitators providing appropriate translation services between the internally used representation formalism and a declarative form). An intrinsically distributed system, WERKL thus can benefit from the results of initiatives such as the ARPA/ISO Intelligent Integration of Information (I^3) Program (e.g. the Agile or Cosmos Projects); the Knowledge Sharing Effort comprising KQML, KIF and Ontolingua; and the Services and Information Management for decision Systems (SIMS) line of research pursued at the ISI. The WERKL architecture lends itself to a federated implementation according to the paradigm of agents, with information agents encapsulating subtasks such as the maintenance of ontologies, the processing of queries across multiple ontologies, and the communication with the user.

4 APPLICATION FIELDS

In a first empirical evaluation, WERKL will be applied in three ongoing projects: MEDCAL, BDB, and HYSAT.

[8] *"In many cases the expressive power of DL formalisms is, strictly speaking insufficient. Nevertheless, sets of DL axioms which "approximately" represent the specified knowledge might be useful. . . . Often, sets of DL axioms which represent more general expressions can be created. These axioms represent the specified knowledge, but due to the generality of the constructions, these axioms can be treated in an unintended way (not corresponding to the knowledge specifications). . . . Note that the usage of generalised representations is dangerous, since inferences which seem obvious to users may not be drawn"* [49].

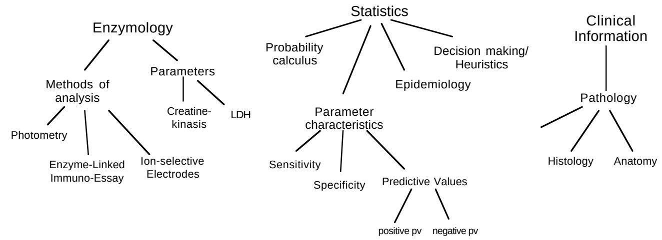

Figure 5. Dimensions involved in clinical decision making

The first of the demonstration applications, MEDCAL is being realized in a medical computer-based learning environment. Here an emphasis is placed on conveying explicitly the interrelationships existing between architectonic spaces that are introduced separately during the medical curriculum, depending on the particular medical task at hand. In the context of teaching the hypothesis-guided decision making process using the assessment of the presence of heart disease as an specific scenario, the way different subject areas (Figure 5) influence the procedure via their respective contributions to the selection of tests and the interpretation of the test results is highlighted. Another functionality to be provided is the dynamic generation of guided tours which fill in the ground between presupposed basic general knowledge about the domain and specific pieces of information meeting particular information needs. These tours shall exemplify how the system can provide guidance during the exploration process when sufficient contextual information is available.

The BDB project is aimed at building multimedia catalogues of products for the building construction industry. In this project, the index layer is used to represent the different views of experts from the building industry. The exploration process allows the discovery of products that users have not been aware of prior to the interaction with the system. It also supports users in tailoring the system according to their needs.

HYSAT (Hypertext System for Architectural Typology [23]) is a project targeted at developing a design-supporting tool for students of architecture. Within this project, multimedia documents are used to represent design examples. Multiple indices focus on the various aspects found in building design (e.g., construction, form, function). Using the system, students learn both about the domain and the structure of the domain space in terms of examples and abstract concepts.

5 RELATED WORK

5.1 Hypermedia

Hypermedia technology has traditionally been considered a "natural" choice for open-world systems [10], even though the level of assistance a computer system can provide is clearly related to the amount of domain knowledge that is explicitly represented within it.[9] Early related research which was aimed at developing systems supporting designers' argumentation activities was originated by Rittel [40] whose work in turn drew on the Issue Based Information Systems (IBIS) introduced by Kunz [26]. Frisse's introduction of *index* and *document spaces* [14] — roughly corresponding to our notions of architectonic and semantic spaces — as complements for navigation is but one example for the subsequent increased research interest in explicit structuring of the knowledge embodied in the semantic

[9] Statistical approaches (e.g., [1]) notwithstanding.

space. Frisse and Cousins also address a number of other important issues, such as human-computer collaboration aspects (*"We believe that most readers will benefit from imprecise results and that they will recognise circumstances where our algorithms are not performing effectively."*), and support for changes of focus.

Initially, the index space was used for the sole purpose of facilitating "global navigation", a role from which it was soon emancipated, as the relevance of navigation and analysis of the index space was recognised [28, 45]. However, later experiences such as the series of research projects leading from SEPIA to DOLPHIN [50, 51, 22, 31] or from Aquanet to HOS and VIKI (a paradigmatic example of spatial hypertext), [32, 33, 34, 47, 35] pointed out a number of serious difficulties related to the introduction of explicit representations of knowledge, the sum of which resulted in users rejecting the formalization of information. These problems include the premature commitment to specific knowledge structures, problems in (re-)use of existing structures by average end-users and the converse problem of defining sharable formalizations, the phenomenon of tacit knowledge, the cognitive overheads involved, and many more (see e.g., [46] for a more detailed overview). Another dissatisfactory line of development is illustrated by the history of the MacWeb system [38, 39] which in the present context can be interpreted as a devolution from an open-world system to an application of descriptive logic to closed-world domains. In contrast, an example of how the insights gained by these experiences can be exploited more fruitfully — albeit at the cost of the restriction to a pure information retrieval system – is given by the recent developments within the MORE multimedia information retrieval project [29, 30].

In WERKL we do not anticipate the adoption of automated feature extraction techniques for media contained in the data layer. Consequently, abstraction and reification of the content has do be done by hand. However, we believe that we can compensate for what at first sight might appear to be a limitation by providing a distributed collaborative environment, where different users are supported in their formalization of parts of the system's content according to their expertise. In addition to the related research just mentioned, the feasibility of our approach is corroborated by other germane work, e.g. in the ComMentor system [41]; the Distributed Link Service derived from the Microcosm project [9]; or the family of ASK systems [42, 43, 2] which has proven successful for settings with static and well-defined user information needs. In comparison to these efforts, the approach taken by WERKL provides a significantly richer index layer in terms of expressiveness of the relations provided to structure the underlying data, which in turn allows for machine support at the knowledge level.

5.2 Case-Based Reasoning

The acknowledgement of the relevance of an appropriate coverage of both architectonic and semantic spaces is further documented in the development history of case-based reasoning systems [53]. The ARCHIE project [17, 12] is characterized by a gradual shift of emphasis towards the semantic space, which went hand in hand with the abandonment of a detailed structuring of the architectonic space. While ARCHIE could be said to suffer from an overly detailed decomposition of the covered design cases, ARCHIE-2 tried to overcome problems by using informal story-based case representations. This group of systems has proven very successful for settings with well-defined user information needs, albeit at the cost of the high human resource demands of the question-based indexing method.

An imbalanced coverage of the architectonic and the semantic spaces thus has been shown to have a strong impact on the usability of a system: users are either confronted with a detailed vocabulary placing a high cognitive burden and imposing artificial restrictions on the accessible domain space, or with brittle systems that fail to provide assistance for unprecedented usage patterns. This is reflected in the recent work on DEDAL and DE-KART [3, 4, 5]. While DEDAL's conceptual index is based on a model of the artifact being designed, a growing number of *proximity retrieval heuristics* concurrently tap into the semantic space and thereby ensure that the users can also access parts of the architectonic space that are not already explicitly incorporated in the system's knowledge base. In this interpretation, although currently restricted to the domain of retrieval of technical documentation, DE-KART can be taken as an example of the kind of incremental extension of the architectonic space we discussed in the previous section.

The aim of the FABEL project is to support design by case-based and model-based methods. These methods are being investigated on the example of designing buildings with complex installations. In its current version, FABEL relies heavily on the KADS methodology: a custom programming language, MoMo, was implemented with the primary goal of providing a means for straightforward generation of executables from KADS specifications. Consequently many of the strengths (e.g., the clear differentiation between different kinds of knowledge — domain and control knowledge — to be considered) but also weaknesses (e.g., the poorly developed strategic knowledge layer was omitted altogether in MoMo) of KADS were inherited [15]. Not too surprisingly, the current emphasis on case-based reasoning methods leads to difficulties related to the open-world domain setting that are comparable to those just described for other systems. For instance, the planning model ARMILLA 5 is an abstraction of real planning processes that was "styled in a way to be systematic, homogeneous and simple", at the cost of omitting common and indispensable features of real planning processes, such as backtracking or flexible task decomposition [44]. Even though the basic paradigm pursued is one of collaborative problem solving between human and computer, at the present stage there exists a very clear-cut distinction between single tasks which can be assigned to system and the majority of others which users have to take on on their own.

6 CONCLUSION

In this paper we have presented a multimedia information system architecture, WERKL, which puts an emphasis on the role of the computer as a partner that augments the capabilities of the human user and offers opportunities to characterize and evaluate the task at hand from different points of view. In this scenario we take on the issues of collaboration, knowledge sharing, and knowledge retrieval in multimedia systems, while trying to maintain a balance between the uses of semantic and architectonic spaces. The WERKL architecture comprises two layers which represent the captured knowledge about the world: a semi-formal hypermedia data layer and a formal index layer. The data layer exploits various media and simple link structures to render aspects of the world that are of interest to the users. The index layer adds meta-knowledge to this structure which enables the machine to draw inferences about the content of the media and its usefulness in regard to given tasks. Building on this basic architecture and the associated exploration cycle as fundamental information gathering mechanism we are now pursuing its implementation following the federated agent architecture model which will be evaluated in the described application fields.

ACKNOWLEDGEMENTS

We would like to thank the referees for their detailed comments which helped improve this paper. The Austrian Research Institute for Artificial Intelligence is supported by the Austrian Federal Ministry for Science, Transport, and the Arts.

REFERENCES

[1] Allan J., Salton G.: The Identification of Text Relations Using Automatic Hypertext Linking, in: Mayfield J. & Nicholas C., *Proceedings Workshop on Intelligent Hypertext, (CIKM 93)*, Arlington, VA, USA, 1993.

[2] Bareiss R., Osgood R.: Applying AI Methods to the Design of Exploratory Hypermedia Systems, in: *Hypertext '93 Proceedings*, ACM, New York, 94–105, 1993.

[3] Baudin C., Gevins J., Baya V., Mabogunje A.: Dedal: Using Domain Concepts to Index Engineering Design Information, in: *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, Lawrence Erlbaum, Hillsdale, NJ, 702–707, 1992.

[4] Baudin C., Underwood J.G., Baya V.: Using Device Models to Facilitate the Retrieval of Multimedia Design Information, in Bajcsy R.(ed.), in: *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, Morgan Kaufmann, San Mateo, CA, 1237–1245, 1993.

[5] Baudin C., Kedar S., Pell B.: Increasing Levels of Assistance in Refinement of Knowledge-based Retrieval Systems, in: *Special Issue: The Integration of Machine Learning and Knowledge Acquisition*, *Knowledge Acquisition*, 6(2), 1994.

[6] Bieber M., Isakowitz T.(eds.): *Communications of the ACM*: Designing Hypermedia Applications, 38(8), 1995.

[7] Borgida A.: On the Relationship between Description Logic and Predicate Logic Queries, in: Adam N.R., et al.(eds.), *Proceedings of the Third International Conference on Information and Knowledge Management* (CIKM94), ACM, New York, 219–225, 1994.

[8] Brachman R.J., Ciccarelli E., Greenfeld N., Yonke M.: KL-ONE Reference Manual, BBN, Cambridge, MA, Report 3848, 1978.

[9] Carr L., Roure D.De, Hall W., Hill G.: The Distributed Link Service: A Tool for Publishers, Authors and Readers, Multimedia Research Group, Department of Electronics & Computer Science, University of Southampton, 1995.

[10] Conklin J.: A Survey of Hypertext, MCC, 1987.

[11] Dillon A., McKnight C., Richardson J.: Space — The Final Chapter or Why Physical Representations are not Semantic Intentions, in: McKnight C., Dillon A., Richardson J. (eds.): *Hypertext: A Psychological Perspective*, Ellis Horwood, 169–191, 1993.

[12] Domeshek E., Kolodner J.: A Case-Based Design Aid for Architecture, in: Gero J.S.(ed.), *Artificial Intelligence in Design 1992*, Kluwer, Boston, 1992.

[13] Fox E.A., Akscyn R.M., Furuta R.K., Leggett J.J.(eds.): *Communications of the ACM*: Digital Libraries, 38(4), 1995.

[14] Frisse M.E., Cousins S.B.: Information Retrieval from Hypertext: Update on the Dynamic Medical Handbook Project, in: *Hypertext'89 Proceedings*, ACM, New York, Special Issue-SIGCHI Bulletin, 199–212, 1989.

[15] Gebhardt F., Bartsch-Spörl B., Oertel W., Walther J.: Anforderungen an die FABEL-Sprache, GMD, Sankt Augustin, Germany, FABEL-Report 11, 1993.

[16] Genesereth M.R., Fikes F.E.: Knowledge Interchange Format, Computer Science Dept., Stanford University, Stanford, CA 94305, USA, 1992.

[17] Goel A.K., Kolodner J.L., Pearce M., Billington R., Zimring C.: Towards a Case -Based Tool for Aiding Conceptual Design Problem Solving, in: *Case-based Reasoning: Proceedings of the 1991 DARPA Workshop*, Morgan Kaufmann, Los Altos, CA, 109–120, 1991.

[18] Gronbaek K., Trigg R.H. (eds.): *Communications of the ACM*: Hypermedia, 37(2), 1994.

[19] Gruber T.R.: A Translation Approach to Portable Ontology Specifications, Knowledge Systems Laboratory, Computer Science Dept., Stanford University, Stanford, CA 94305, USA, Knowledge Acquisition 5(2):199–220, 1992.

[20] Gruber T.R.: Ontolingua: A Mechanism to Support Portable Ontologies, Knowledge Systems Laboratory, Computer Science Dept., Stanford University, Stanford, CA 94305, USA, 1993.

[21] Guarino N., Giaretta P.: Ontologies and Knowledge Bases, in: Mars N.J.I. (ed.): *Towards Very Large Knowledge Bases*, IOS Press, Amsterdam, 1995, 25–32, 1995.

[22] Haake J.M., Neuwirth C.M., Streitz N.A.: Coexistence and Transformation of Informal and Formal Structures: Requirements for More Flexible Hypermedia Systems, in: *ECHT'94 Proceedings*, ACM, New York, 1–12, 1994.

[23] Herzog M.: *The Use of Intelligent Hypermedia in Architectural Design Environments — a Conceptual Framework*, Institut für Informationssysteme, Abteilung für Datenbanken und Expertensysteme, Technische Universität Wien, Diploma Thesis, 1994.

[24] Kaplan N., Moulthrop S.: Where No Mind Has Gone Before: Ontological Design for Virtual Spaces, in: *ECHT'94 Proceedings*, ACM, New York, 206–216, 1994.

[25] Kashyap V., Shah K., Sheth A.: Metadata for Building the MultiMedia Patch Quilt, in Subrahmanian V.S. & Jajodia S.(eds.), Multimedia Database Systems, Springer, Berlin, 297–319, 1996.

[26] Kunz W., Rittel H.: Issues as Elements of Infromation Systems, Institut für Grundlagen der Planung, Stuttgart, Working Paper 131, 1970.

[27] Lambrix P., Padgham L.: Part-of Reasoning in Description Logics: A Document Management Application, in Borgida A., et al.(eds.), *International Workshop on Description Logics*, Rome, June 2–3, Dipartimento di Informatica e Sistemistica, Universita' di Roma, “La Sapienza”, Italia, 106–108, 1995.

[28] Lucarella D.: A model for hypertext-based information retrieval, in Rizk A.(ed.), *Proceedings of the First European Conference on Hypertext*, Cambridge University Press, New York, The Cambridge Series on Electronic Publishing, 1990.

[29] Lucarella D., Parisotto S., Zanzi A.: MORE: Multimedia Object Retrieval Environment, in: *Hypertext '93 Proceedings*, ACM, New York, 39–50, 1993.

[30] Lucarella D., Zanzi A.: A Visual Retrieval Environment for Hypermedia Information Systems, *ACM Transactions on Information Systems*, 14 (1), 3–29, 1996.

[31] Mark G., Haake J.M., Streitz N.A.: The Use of Hypermedia in Group Problem Solving: An Evaluation of the DOLPHIN Electronic Meeting Room Environment, in: *Proceedings of European Conference on Computer-Supported Cooperative Work (E-CSCW'95)*, Stockholm, September 10-15, 1995.

[32] Marshall C.C., Halasz F.G., Rogers R.A., Janssen W.C.Jr: Aquanet: a hypertext tool to hold your knowledge in place, in: *Third ACM Conference on Hypertext (Hypertext '91): Proceedings*, ACM, New York, 261–275, 1991.

[33] Marshall C.C., Rogers R.A.: Two Years before the Mist: Experiences with Aquanet, in Lucarella D., et al. (eds.), *Fourth ACM Conference on Hypertext*, ACM Press, New York, 53–62, 1992.

[34] Marshall C.C., Shipman F.M.III: Searching for the Missing Link: Discovering Implicit Structure in Spatial Hypertext, in: *Hypertext '93 Proceedings*, ACM, New York, 217–230, 1993.

[35] Marshall C.C., Shipman F.M.III, Coombs J.H.: VIKI: Spatial Hypertext Supporting Emergent Structure, in: *ECHT'94 Proceedings*, ACM, New York, 13–23, 1994.

[36] Mili H., Rada R.: Medical expertext as regularity in semantic nets, in: *Artificial Intelligence in Medicine*, 2(4), 1990.

[37] Muggleton S., Srinivasan A., Bain M.: *MDL Codes for Non-Monotonic Learning*, Turing Institute, Glasgow, 1991.

[38] Nanard J., Nanard M.: Using Structured Types to Incorporate Knowledge in Hypertext, in *Third ACM Conference on Hypertext (Hypertext '91): Proceedings*, ACM, New York, 329–344, 1991.

[39] Nanard J., Nanard M.: Should Anchors Be Typed Too? An Experiment

with MacWeb, in *Hypertext '93 Proceedings*, ACM, New York, 51–62, 1993.

[40] Rittel H.: On the Planning Crisis: Systems Analysis of the 'First and Second Generations', *Bedriftsokonomen*, No. 8, 390–396, 1972.

[41] Roescheisen M., Mogensen C., Winograd T.: A Platform for Third-Party Value- Added Information Providers: Architecture, Protocols, and Usage Examples, Stanford Integrated Digital Library Project, Computer Science Dept., Stanford University, 1994, updated April, 1995.

[42] Schank R.: Rules and Topics in Conversation, in: *Cognitive Science*, Vol. 1, 421–441, 1977.

[43] Schank R., Cleary C.: Engines for Education, Lawrence Erlbaum Ass., New Haven, 1994.

[44] Schmidt-Belz B.: Scenario of FABEL Prototype 3 Supporting Architectural Design, GMD, Sankt Augustin, Germany, FABEL-Report 40, 1995.

[45] Schnase J.L., Leggett J.J., Hicks D.L., Szabo R.L.: Semantic Data Modeling of Hypermedia Associations, Hypermedia Research Laboratory, Department of Computer Science, Texas A&M University, College Station, TX, TAMU-HRL-91-005, 1991. Also in: *ACM Transactions on Information Systems*, 11(1), 27–50, 1993.

[46] Shipman F.M.III: Supporting Knowledge-Base Evolution with Incremental Formalization, University of Colorado at Boulder, Department of Computer Science, 1993.

[47] Shipman F.M.III, Marshall C.C.: Formality Considered Harmful: Experiences, Emerging Themes, and Directions, Xerox Palo Alto Research Center, ISTL-CSA-94-08-02,, 1994.

[48] Smithers T. and Troxell W.: Design is intelligent behaviour, but what's the formalism?, in: *AI EDAM*, 4(2):89–98, 1990.

[49] Speel P.-H.: Selecting description logics for real applications, in Borgida A., et al.(eds.), *International Workshop on Description Logics*, Rome, June 2– 3, Dipartimento di Informatica e Sistemistica, Universita' di Roma, "La Sapienza", Italia, 143–147, 1995.

[50] Streitz N.A., Hannemann J., Thuering M.: From Ideas and Arguments to Hyperdocuments: Travelling Through Activity Spaces, in: *Hypertext'89 Proceedings*, ACM, New York, Special Issue-SIGCHI Bulletin, 343–364, 1989.

[51] Streitz N., Haake J., Hannemann J., Lemke A., Schuler W., Schuett H., Thuering M.: SEPIA: A Cooperative Hypermedia Authoring Environment, in Lucarella D., et al.(eds.), *Fourth ACM Conference on Hypertext*, ACM Press, New York, 11-22, 1992.

[52] Terveen L.G.: Overview of Human-Computer Collaboration, in Special Issue: Human-Computer Collaboration, *Knowledge-Based Systems*, 8(2/3), 67–81, 1995.

[53] Watson I., Marir F.: Case-Based Reasoning: A Review, The Knowledge Engineering Review, 9(4), 327–354, 1994.

[54] Wexelblat A.: Giving Meaning to Place: Semantic Spaces, in: Benedikt M.(ed.), *Cyberspace — First Steps*, MIT Press, Cambridge, MA, 255–272, 1991.

[55] Wolverton M., Hayes-Roth B.: *Finding Analogues for Innovative Design*, Stanford University, Knowledge Systems Laboratory, KSL-95-32, 1995.