Subject-Oriented Work: Lessons Learned from an Interdisciplinary Content Management Project

Joachim W. Schmidt, Hans-Werner Sehring, Michael Skusa, and Axel Wienberg

Technical University TUHH Hamburg, Software Systems Institute,
Harburger Schloßstraße 20, D-21073 Hamburg, Germany
{j.w.schmidt, hw.sehring, skusa, ax.wienberg}@tuhh.de

Abstract. The two broad cases, data- and content-based applications, differ substantially in the fact that data case applications are abstracted first before they cross any system boundary while for content cases it is the system itself which has to map application content into some data-based technology. Through application analysis and software design we are aware of the difficulties of such mappings. In an interdisciplinary project with our Art History colleagues who are working in the subject area of "Political Iconography" we are gaining substantial insight into their *Subject-Oriented Working* (SOWing) needs and into initial requirements for a SOWing environment. In this paper we outline the project, its basic models, their generalization as well as our initial experiences with prototypical SOWing implementations. We emphasizes the conceptual and terminological aspects of our approach, sketch some of the technical requirements of a generic SOWing software platform and relate our work to various XML-based activities.

1 Introduction

As a result of advanced and extensible database technology now being available as off-the-shelf products, a substantial part of database research and development work has generalized into work on models and systems for multimedia content management. R&D in content management includes a range of models and systems concentrating on services for the following three lines of work:

- content production and publication work using multimedia documents;
- classification and retrieval work based on document content;
- management and control of such work for communities of users differentiated by their roles and rights, interests and profiles, etc.

The work reported in this paper is based on an interdisciplinary project with a partner from the humanities with strong semantic and weak formal commitment. Our project partner specializes on work in icon- and text-based content from the subject area of "Political Iconography". This content is organized as a paper- and drawer-based subject index (PI-"Bildindex", BPI, see fig. 1) and is used for Art History research and education [33].

Iconographic work has a long tradition in Art History, dating back to 19th century "Christian Iconography", and is based on integrated experience from three sources:
- Art: multimedia content;
- Art History: process knowledge;
- Library Sciences: subject-oriented content classification and retrieval.

In our context we use the notion of *subject-orientation* very much in the sense of library science, as, for example, stated by Elaine Svenonius: "In a subject language the extension of a term is the class of all documents about what the term denotes, such as all documents about butterflies." This understanding differs substantially from natural language where "the extension, or extensional meaning, of a word is the class of entities denoted by that word, such as the class consisting of all butterflies" [30]. And both understandings are in clear contrast to the semantics of terms in programming languages and database models.

The interdisciplinary project "Warburg Electronic Library (WEL)" models and computerizes BPI content and services [3], [19], and the WEL prototype allows interdisciplinary experiments and insights into multimedia content management and applications.

The overall goal of the WEL project and the follow-up SOWing project is
- the generalization of our subject-oriented working experience,
- a work plan for R&D in subject-oriented work and content management, and
- a generic Subject-Oriented Working environment (SOWing environment).

Currently, many contributions to such R&D are based on XML as a syntactic framework which provides a structural basis as well as some form of implementation platform. The main reasons for XML's powerful position are its strong structural commitment and its semantic neutrality.

Successful content management requires that the three lines of work
- content production and publication work by multimedia documents;
- classification and retrieval work based on document content;
- management and control of such work for communities of users differentiated by their roles an rights, profiles and interests, etc. [21]

are not supported in isolation but in a coherent and cooperative working environment covering the entire space spanned by all three dimensions.

The main reason why XML-based work on content management often falls short can be stated as a corollary of XML's strength (see above): its weak semantic and exclusively structural commitment. Much of the XML-based R&D contributes to the above three lines of work only individually. Examples include [7], [31].

The paper is structured as follows. Section two introduces the two projects involved, the "Bildindex für Politische Ikonographie (BPI)" and the "Warburg Electronic Library (WEL)". In section three the WEL model is generalized towards a generic "Work Explication Language". A system's view of the WEL prototype is described in section four and the first contours of a generic Subject-Oriented Working environment (SOWing platform) are outlined. Related work, in particular work in the XML context, is discussed in section five. The paper concludes with a short summary and a reference to future work in our long-term SOWing project.

2 Warburg Electronic Library: An Interdisciplinary Content Management Project

The development of the currently predominant data management models was heavily influenced by application requirements from the business and banking world and their bookkeeping experience: the concepts of record, tabular view, transaction etc. are obvious examples. Data model development had to go through several generations – record-based file management, hierarchical databases and network models – until the relational data model reached a widely accepted level of abstraction for database structuring and content-based data operation.

For traditional relational data management we basically assume that content is "values of quantified variables" from business domains operated by transactions and laid out as tables with rows and columns. The questions arise:

- How can we generalize from data management to the area of content management where content domains are not "just dates and dollars", content operation goes beyond "debit-credit transactions" and content layout means multimedia documents?
- What are key application areas beyond bookkeeping which help us understand, conceptualize and finally implement the core set of requirements for multimedia content management in terms of domain modelling, content-oriented work support as well as content (re-) presentation?

Fig. 1. Working with the Index for Political Iconography in the Warburg-Haus, Hamburg

The work presented here is based on an interdisciplinary R&D-project between Computer Science and Art History, the "Warburg Electronic Library" project. The application area was chosen because of Art History's long-term working experience with content of various media. The project itself is founded on extensive material and user experience from the area of "Political Iconography".

2.1 Subject-Oriented Work in Political Iconography

Political iconography basically intends to capture the semantics of key concepts of the political realm under the assumption that political goals, roles, values, means etc. requires mass communication which is implemented by the iconographic use of images. Our partner project in Art History, the “Bildindex zur Politischen Iconographie (BPI)”, was initiated in 1982 by the Art Historian Martin Warnke [33] and consists of roughly 1,500 named political concepts (*subject terms*, “Schlagworte”) and more than 300,000 records on iconographic works relevant to the BPI. In 1990 Warnke’s work was awarded the Leibniz-Preis, one of the most prestigious research grants in Germany.

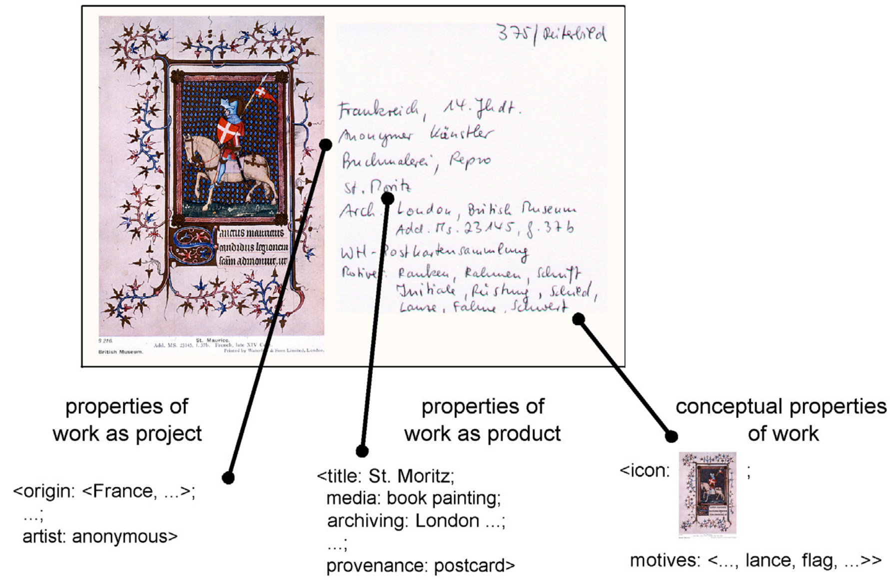

Fig. 2. BPI “Bildkarte” St. Moritz (image card) describing art work by attribute aggregation

Starting with this experience, BPI work essentially relies on an Art Historian’s knowledge of (documents referring to) political acts in which images play an active role. Art Historians interpret “acts” as encompassing aspects of

- “projects” (who initiated and contributed to an act? the when and where of an act? etc.);
- “products” (what piece of art did the project produce? on what medium? place of current residence etc.); and finally, the
- “concepts” behind the act (what political goals, roles, institutions etc. are addressed? what iconographic means are used by the artist? etc.).

On this knowledge level, BPI work identifies political concepts and names them individually by *subject terms* – e.g., by “ruler”, “prince”, “pope”, “equestrian statue”, etc.

Subject term semantics is methodologically captured and systematically represented in the BPI by the following steps:

1. designing a conceptual, prototypical and representative (mostly mental) model for each subject term, e.g., a prototypical equestrian statute [29] ;
2. giving value to the relevant variables or facets of such prototypes by reference to the Art Historian's knowledge of "good cases", i.e., political acts with an iconographic dimension. Each such variable or facet is represented by a BPI entry ("Bildkarte", "Textkarte", "Videokarte" – "media card" etc.) which holds a description of a "good case" for that facet, see for example, St. Moritz, fig. 2.
3. collecting all BPI entries on the same prototype into a single extent ("Bildkartenstapel", …, "stack of media cards", see fig. 3) thus defining the semantics of a subject term. Additional fine structure may be imposed on subject term extents (order, "neighborhood", named subextents, general association/navigation etc.);
4. maintaining a ("completion") process aiming at a "best possible" definition of the subject area at hand by
    - "representative" subject terms covering the subject area at hand;
    - "qualifying" prototypes for each subject term;
    - "complete" sets of facets for prototype description;
    - "good" cases for facet substantiation.

This makes it quite clear that the BPI is by no means just an index for accessing an image repository. The BPI uses images only in their rather specific role as icons and for the specific purpose of contributing to the description of cases and thus to the semantics of subject terms [32]. In this sense, images represent the iconographic vocabulary of BPI documents just as keywords contribute to the linguistic vocabulary of text documents.

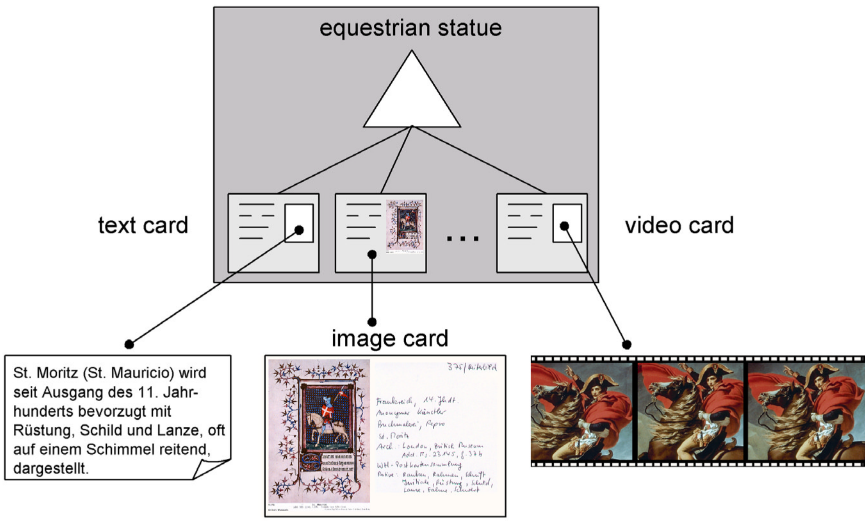

Fig. 3. BPI subject term semantics (e.g. equestrian statue) by media card classification (image, text, video cards etc.)

Art Historians with their long tradition of working with content represented by multiple media are far from restricting themselves to a mainly technical view on multimedia as most of the currently booming projects in Computer Science seem to do. Our Art History colleagues are much closer to the message of people such as Marshall McLuhan who understand media as “extensions of men [18]”.

The BPI has essentially two groups of users:
- a few highly experienced BPI editors for content maintenance and
- various broader user communities which access BPI content for research and education purposes.

Being implemented on paper technology, the traditional BPI shows severe conceptual and technical shortcomings:
- conceptually: the above attributes “representative”, “good”, “complete”, etc. are highly subjective and, therefore, “completion semantics” is hard to meet even within a “single-person-owned” subject index;
- technically: severe representational limitations are obvious and range from single subsumption of BPI entries to a lack of online and networked BPI access.

In the subsequent section we outline two contributions of the “Warburg Electronic Library” project which approaches the above conceptual and technical shortcomings through an advanced Digital Library project which, as a prime application, is now hosting the “Index for Political Iconography”.

2.2 A Subject-oriented Working Environment: Warburg Electronic Library

Viewed from our Computer Science perspective which shifted in recent years from basic research in “persistent database programming” towards R&D in “software systems for content management (online, multimedia, …)”, the WEL project addresses a range of highly relevant and interrelated content application issues:
- content representation by multiple media: images, texts, data, …;
- content structuring, navigation and querying; content presentation;
- content work exploiting subjects and ontologies: classification, indexing, …;
- utilization of different referencing mechanisms: icon, index, symbol;
- cooperative projects on multimedia content in research and education.

The WEL is an interdisciplinary project between the Art History department of Hamburg University (Research Group on Political Iconography, Warburg-Haus, Hamburg) and the Software Systems Institute of the Technical University, TUHH, Hamburg. It began in 1996 as a 5-year project and will be extended into an interdisciplinary R&D-framework involving several Hamburg-based institutions.

For a short WEL overview we will concentrate on two project contributions:
- semantic modelling principles for WEL-design;
- personalized digital WEL libraries based on project-specific prototypes and their use in Art History education.

WEL Semantic Modelling Principles. The WEL design is based – as is already the BPI design – on the classical semantic data modelling principles [28], [6]: aggregation, classification, generalization / specialization and association / navigation (see figs. 2 and 3).

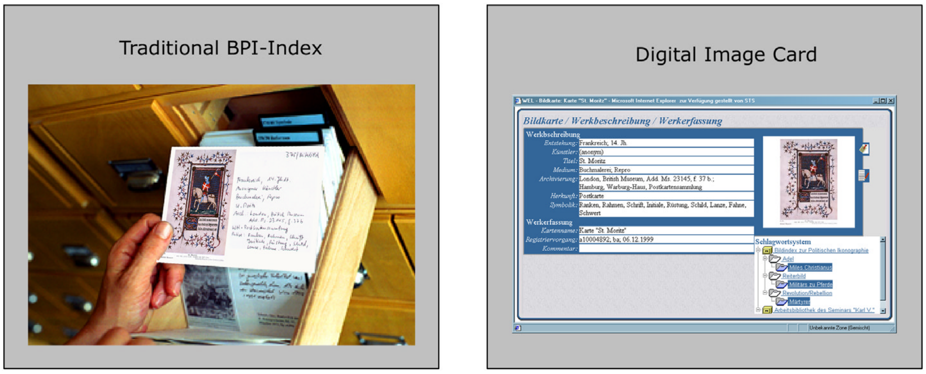

Fig. 4. Media card associated with (multiple) subject terms and with information on classification work

However, it is important to note that the semantics of subject classes and their entries originate from different semantic sources and, therefore, go beyond classical data modelling (see also sec. 4.2):

- *object semantics*: seen from a data modelling point of view, subject class entries are also entities of some object classes in the sense of object-oriented modelling. However, a subject class extent may be heterogeneous because its entries may describe documents of different media – texts, images, videos etc. Therefore, subject class entries viewed as objects may belong to different object classes – text, image, video classes etc.
- *content semantics*: furthermore, all content described by the entries of the same subject class shares some semantic key elements. All BPI documents referring, for example, to the subject class “ruler” make use of graphical textual key icons such as swords, crowns, scepters, horses etc.; similarly, the text documents associated with a certain subject class contain overlapping sets of subject-related keywords. Such sets of key icons capture essential parts of content and, thus, of subject term semantics. Note that specialization of subject classes goes along with extension and union of subject-related key icon sets while generalization relies on reduction and intersection.
- *completion semantics*: in section 2.1 we referred to the soft semantic constraint of achieving best possible subject definition as “completion semantics”. Although this may be considered more as an issue of class *pragmatics* than class semantics it implies formal constraints on subject class extents. Since users of subject definitions act under the assumption that the subject owner established a subject extent which represents *all* relevant aspects of the owners subject prototype, any change of that extent is primarily monotonic, i.e., extents of subject terms are only changed by

adding or replacing its entries. Therefore, references to subject class entries should not become invalid.

Figure 4 shows a media card for St. Moritz together with the (multiple) subject terms to which St. Moritz contributes. The example also links St. Moritz to details of the classification process by which this card entered the WEL. This information is essential for the realization of project-oriented views on subject terms, or reference libraries:

- thematic views (*customization*): projects usually concentrate on sub-areas of the all encompassing "Index for the Political Iconography";
- personalized views (*personalization*): they cope with the conceptual problems with "completion semantics" mentioned above.

Initial experiences with both of these viewing mechanisms are outlined in the subsequent section.

Subject-Oriented Work in Art History Education. A key experience of the WEL project relates to the two dimensions of subject-oriented work: subject-orientation as a thematic view (*customization*) and as an individualized view (*personalization*). Speaking in terms of Digital Libraries both dimensions are approached by the WEL concept of "reference libraries" ("Handbibliothek"), which are essentially SOWing environments customized and personalized according to the requirements of individual projects or persons [17].

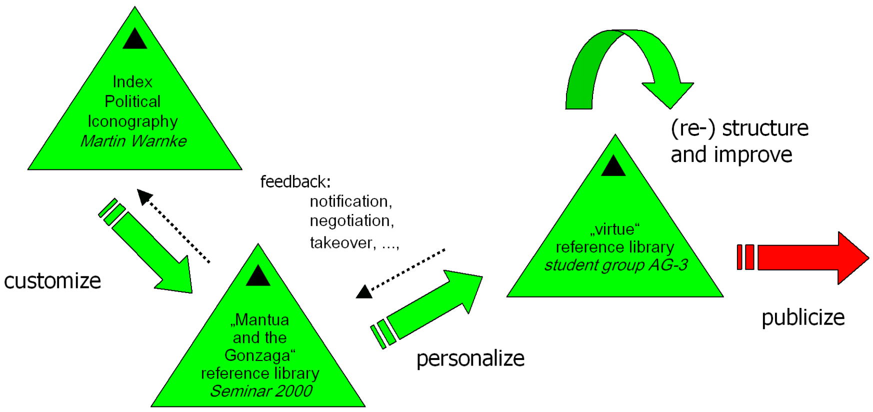

Fig. 5. Thematic and personalized subject views for the "Mantua" seminar

Figure 5 outlines the use of customized and personalized SOWing environments in an Art History seminar on "Mantua and the Gonzaga" [26]. The general BPI ("Warnke-owned") is first customized into a subject index for the "Mantua and the Gonzaga" seminar project. The main objective of the individual student projects in that seminar is to further personalize the seminar index, structurally and content-wise, and to produce, for example, a project-specific subject index for topics such as "Studiolo" or "Camera picta" [20]. Publicizing the final subject content in some form of media document (see fig.6) - traditional print report or interactive website - constitutes another seminar objective [34].

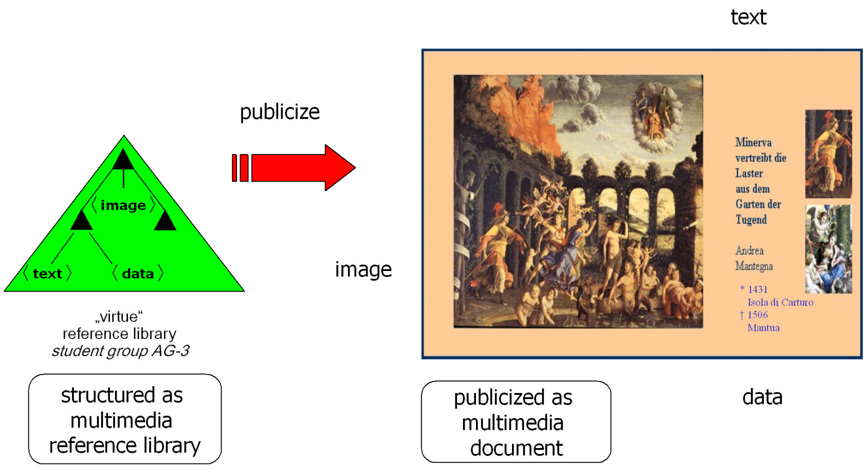

Fig. 6. The "virtue" reference library publicized as a multimedia document

3 Towards a Generalized WEL-Model for Subject-Oriented Work

The prime experience gained from our interdisciplinary WEL project is a deeper insight into the intertwining of SOWing entities and their working relationships. Conceptually, the services of our SOWing platform are primarily based on four kinds of entities which are straight-forward generalizations of the corresponding WEL entities:
- "work cases" as the basic abstraction of the acts and entities of interest in a domain;
- "case documents" which are abstract or physical entities reporting on such work;
- "case entries" which record the essence of work case documents;
- "subject terms" which, based on such case entries, structure the domain, define its semantics, and give access to its documents and works.

Since all four notions are quite generic, there is ample space for generalizing in our SOWing project the models and systems for subject-oriented work far beyond our initial WEL approach.

In the subsequent sections we will outline an extended SOWing model and platform and re-interpret the acronym WEL from "Warburg Electronic Library" to "Work Explication Language", very much in the sense of the definition of ontology as "a theory regarding entities, especially abstract entities to be admitted into a language of description" [35].

Subsequently, we discuss the four kinds of SOWing entities in more detail.

3.1 On a Generic Notion of "Work"

Central to the content management support provided by our SOWing platform is a generic notion of "work". Our work concept is based on the WEL experience and is characterized and modeled by three groups of properties (see also fig. 2):
- work as a "project", i.e., the circumstances under which work is performed;
- work as a "product", i.e., a work's result; and
- work as a "concept", i.e., the conceptual idea behind a work.

Figure 7 relates these three work characteristics using a "work triangle" diagram. The generalization of WEL work examples such as the one given by fig. 2 is obvious.

The upper part of fig. 8 depicts a second work case, similarly structured but rather different in nature. Fig. 8.2 depicts a WEL work case done by a Mr. B. when producing a work document description of the Gonzaga-Mantegna-Minerva work case and entering it into the WEL. While there may be only partial knowledge on the renaissance work case – essentially only Mantegna's picture survived – the WEL work case being supported and observed by the SOWing platform can receive an arbitrarily extensive SOWing coverage.

This reflective capability is probably the most powerful and unique aspect of our SOWing approach. Reflection provides the basis for a wide range of services for customization and personalization, self-description and profiling and for all kinds of guiding and tracing support [27].

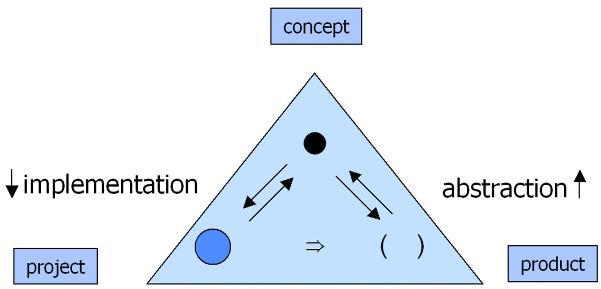

Fig. 7. WEL work structure (work triangle)

Examples of WEL work are presented in fig. 8. The lower part (fig. 8.1) models a specific work case by which a member of the Gonzaga family residing in Mantua during the 15th century asked the artist Mantegna for a painting addressing the issue of virtues and sins. Mantegna chose the goddess Minerva as the central motive and presented her expelling the sin out of the garden of virtue. This 15th century work case may be reported by some publicized work document, most probably, however, the case is just part of an Art Historian's body of knowledge about the Italian renaissance.

3.2 Work Case Documents

Work such as the Gonzaga-Mantegna-Minerva case of fig. 8 is documented typically in narrative form and presented by some multi-media documents – texts, images, speech etc. and combinations thereof. In our SOWing approach such documents are assumed to represent content in terms of the above three dimensions: work project, work product and the concepts behind both (see also fig. 7). This view on documents is quite general and allows interpretation ranging from the rather informal but very expressive documents of "Political Iconography" to partially formalized diagrams such as flow charts and UML entities and to computer programs and operator instructions with a fully formalized semantics.

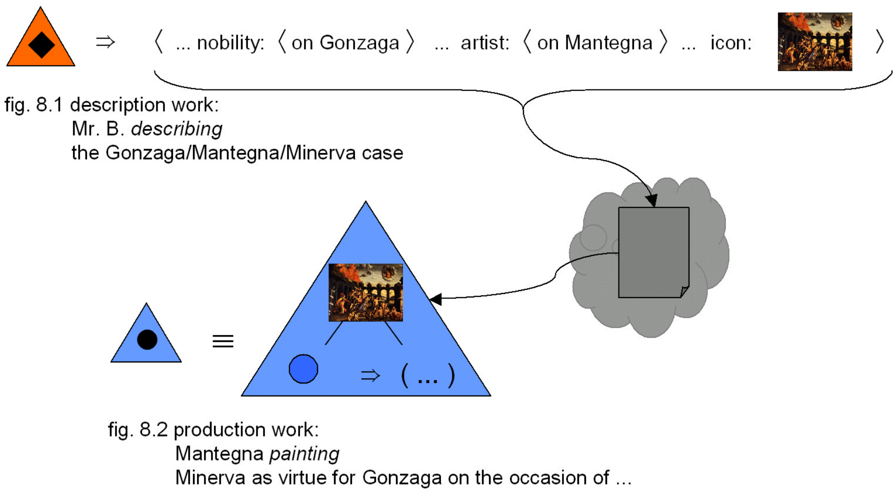

Fig. 8. WEL work graph: production work (fig. 8.1) and description work (fig. 8.2)

Work documents may also vary in terms of their completeness in the sense that the work they report may be known only partially, for example, by its product. Many examples can be found in Art History where most of the project knowledge is usually lost and only the image survives (the opposite case also exists). Nevertheless, we agree with our colleagues from Art History that products should never be considered in isolation but always be recorded in the context of the project (persons, time, place, tools, etc.) for which they were created. For documents produced within a computerized environment this has partially become standard although there is no overall concept of what to do with this information.

Note the dialectic character of our SOWing position in this point: on the one hand side, work cases are assumed to be reported by documents, on the other hand all documents – at least the ones produced by the SOWing environment and its tools – are considered as work cases and, yet again, reported by work case documents and entries.

3.3 Work Case Entries

In the SOWing approach work case documents are described and recorded by work case entries which establish the relationship between subject terms and their semantics on one side and the work cases and their documentation on the other. Such case entries generalize the media cards of the WEL system.

Conceptually there is a third relationship involved which associates a work case entry with a class of icons considered "representative" for the kind of work which is described by the case. We use the term "icon" here in the sense of iconic signifier (as opposed to indexical or symbolic signifiers, see, for example, [9]). For image documents icons specialize to iconographic signifiers, for text documents to keywords ("Stichworte") etc.

A case entry on an image and text document with content on the subject term "ruler", for example, will for its image part be described by characteristic icons from a subject-specific icon set {crown, sword, scepter, ...}, for its textual part by a corresponding set of keywords.

Seen from an *object* point of view, case entries are instances of media-specific object classes (image, text, video classes etc.) while from a *content* perspective they draw constraints from (hierarchies of) icon classes. Finally, from a *subject*-oriented position, case entries become members of the extent of some subject classes (our old WEL stack of media cards) thus contributing to the definition of their semantics.

In section four we will relate and discuss these three perspectives in terms of class diagrams (fig. 10).

3.4 Subject Definition Work

Subject term semantics is essentially defined extensionally by document descriptions. Intentionally their semantics is captured in part by icon classes shared by such extents. Both extensional and intensional semantics are related by the fact that each entry into the subject term's extent shows a characteristic profile over the icon class related to its subject term. A specific image document contributing to the semantics of "ruler" will not display the entire icon class for rulers, i.e., the set {crown, sword, scepter, ...} but a characteristic subset of it, probably a crown and a sword in a prominent position within the image.

Most of the production work on which the BPI is based typically took place outside a computerized environment – usually the referenced iconographic work dates back several centuries. Production work may, however, also mean the production of documents *about* the original iconographic work and such document work may well contribute to or profit from a SOWing environment.

Description support being definitely a matter of the SOWing platform may, for example, provide reference to other subject terms covered by the index [22]. For the Gonzaga-Mantegna-Minerva case it may be quite enlightening to capture the reason *why* Gonzaga ordered that picture by referring to some subject term "virtue" or its generalization "political objectives" to which Machiavelli's work "Il Principe", a successful handbook for renaissance rulers, contributed.

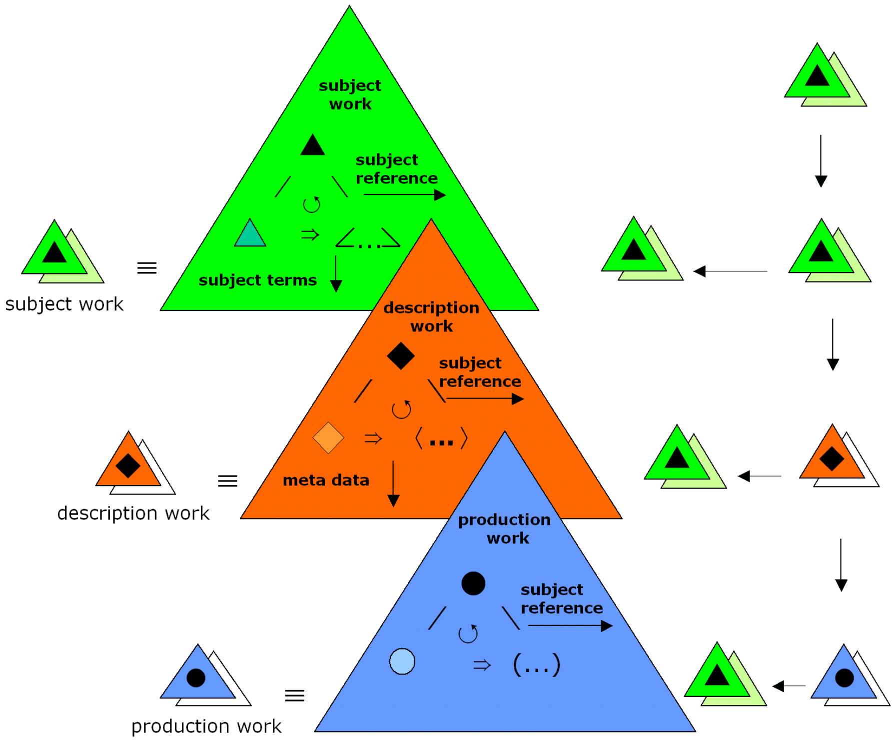

Fig. 9. An overview of Subject-Oriented Work

Fig. 9 gives an overview of the subject-oriented work and its support through a SOWing platform. It relates production and description work to subject work.

As mentioned above, a major group of SOWing services is based on the fact that "work is a first class citizen" in the SOWing world. This implies that work while being supported by the system is automatically identified, described and associated with work-related subject terms (project time and participants, product media and archives etc.). Evaluating such terms provides the basis for substantially improved work support, e.g., work session management, work distribution, protection, personalization etc.

4 The SOWing System

Our field studies provided us with a good basis for user requirements analysis of our SOWing approach and our extensive prototyping experience allowed us us a deeper insight into the architectural and functional alternatives of SOWing system design and implementation.

4.1 WEL Prototype Experience

The WEL prototypes developed so far are the major source of experience on which the SOWing approach is based. The initial version was based on the Tycoon-2 persistent programming environment, an object-oriented, higher-order programming language with orthogonal persistence [16], [23]. Using a Tycoon-based acquisition and cataloging tool the Computer Scientists, Art Historians and many students from both departments digitized many thousands of images and transformed file cards from their physical representation into a digital one. This work is still in progress. The descriptive data in the card catalogue were entered and revised by Art Historians via a web-based editor.

With the evolution of Java from a small object-oriented programming language for embedded devices to a mainstream programming language for networked applications the WEL system was moved to Java-related technology.

The current WEL version is based on an early version of a commercial Content Management System (CMS) [5] which is entirely Java-based and resides on top of a relational database management system.

For our current prototype we were particularly interested in understanding to what extent features of commercial CMS technology meet the requirements of our SOWing platform. It turned out that although the CMS provided several abstractions useful for our SOWing data model important relationships between and inside the subject classes could not be treated as first class objects. Furthermore, the modeling facilities of commercial CMSs are not yet rich enough to meet the needs of our generic SOWing approach. Commercial CMS technology concentrates on specific application domains with editorial processes for more or less isolated units of work (e.g., news articles) with little association to other entities. SOWing requires, however, in addition to storing and retrieving the document itself an extended functionality for the embedding of documents into the SOWing context with all its relationships.

Commercial CMS and DBMS technology supports only the specific set of navigation methods predominant in their main application domain, and our prototype ran into performance problems as soon as users left those default navigation paths. As a consequence, we met the domain specific access requirements by introducing an additional application layer on top of the CMS.

Several graphical web-based user interfaces as well as GUI editors were developed to enable various classes of users to browse through the subject index, create personal indices and collect references to documents accessible via the global index.

In parallel to the Art History project the generic WEL system was adopted to maintain an index created for the management of concepts relevant in advertising. This was and still is carried out in cooperation with a commercial agency from the advertising industry.

4.2 Subject Classes, Object Classes, and Icon Classes

In section three we introduced the four kinds of SOWing entities: work cases, case documents, case entries, and subject terms. While the former two (sections 3.1 and

3.2) are highly relevant for the conceptual foundation of our SOWing approach, the latter two (sections 3.3 and 3.4) are also important for system implementation.

Case entries are related to subject terms by classification relationships. Subject terms are structured in a hierarchy – the subject index – made up by subject term generalization and specialization.

As described in section 2.1 the relationship between case entries and subject terms has a double meaning:

- one the one hand, documents are classified by binding their case entries to subject terms;
- on the other hand, each document contributes to the subject term's definition.

The set of entries chosen to define a subject is supposed to be minimal; only entries which introduce a new and relevant facet of the prototypical extent are added. In this way subject indices reflect the domain knowledge from the perspective of the owner of the subject index.

Subject indices are not only used to capture "primary knowledge" of domains but also "secondary knowledge" as, for example,

- on the organizational and history of a community, including knowledge of users, rights granted to them, etc. [8]
- on the layout and handling of documents, i.e., properties not directly related to their content, e.g., (kind of) origin, document types, quality, etc.

This leads to different types of subject indices some of which are used by the SOWing system itself, e.g., user classification to handle project-specific access rights.

Since entries can contribute to more than one subject term definition they may belong to more than one subject extent maybe in different indices of the same or different types. The semantics of multiple subsumption varies in each of the cases.

Several contributions define the semantics of the instantiation relationship:

- object semantics: technically, descriptions have a type (in the sense of a data type);
- content semantics: in addition to the attributes with object semantics, the use of icons for content description determines a semantic type; the more special a subject becomes the more icons the entries in its extent are expected to have;
- completion semantics: the extent of a subject term is supposed to fully describe its semantics (at the current time, for the user who created it).

It is interesting to see how in our design of SOWing entities traditional elements of object-oriented classes coexist with novel aspects of subject-oriented modelling. This coexistence of object- and subject-oriented semantic elements is illustrated by the diagram in fig. 10 which is based on an extended UML notation.

The upper right third of the diagram shows a traditional class hierarchy of an object-oriented model. On top is the class "Object" representing the root of the class hierarchy. From this a class "Case Entry" is derived which has attributes for the features which all kinds of case entries share. Subclasses of "Case Entry" are introduced for each media type. These classes might introduce further features, e.g., painter for images, author for texts. Instances ("entry" in the diagram) of such classes are constructed in the usual object-oriented manner: the object is created for a given class, so that its structure is known for its whole lifetime.

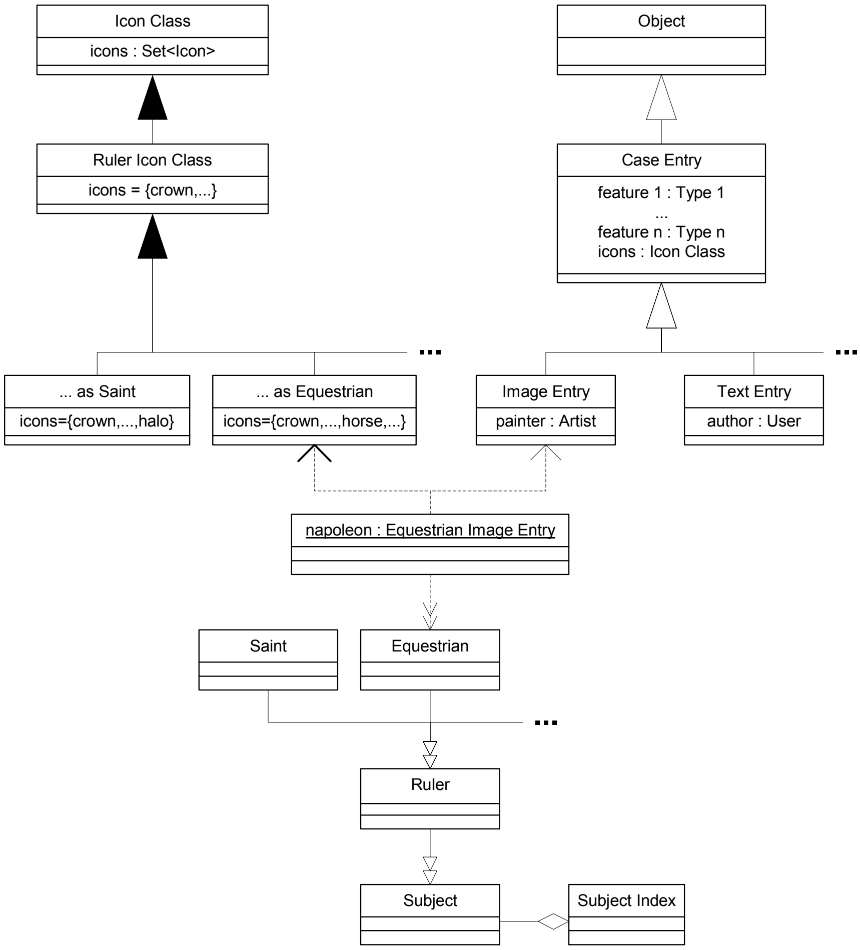

Fig. 10. Subject, object and icon classes

Important for the SOWing model is the special attribute “icons” which is defined by “Case Entry” and which represents references to an icon class. Icon classes describe sets of icons which relate to subject terms. Icons reflect the content of the described document, e.g. keywords in a text or symbols in an image. As indicated in the upper left of the diagram (fig. 10), icon classes are ordered in a specialization relationship which is derived from the inclusion of the icon sets. We use the filled-in arrow head to visualize icon class specialization and a thick arrow head for icon class “instantiation”.

Finally, the lower part of fig. 10 shows a (UML-inspired) formulation of subject classes. We use double-headed arrows for specialization and instantiation. Dual to the (object) class “Object” we introduce a (subject) class “Subject” as the root of the subject class hierarchy. Although subject terms and icon classes are only loosely related, it is very likely that all the entries of the same subject class will have a similar profile over the icons of the corresponding icon class.

The "Subject Index" to which a subject belongs defines the domain to which a subject term contributes. This models the perspectives under which the classification of a description can be viewed. For any given application often only one index will be considered at a time. However, personalization and customization will require the SOWing system to cope internally with several indices simultaneously.

There are two fundamental uses of the structure shown in the diagram, fig. 10:

- if a subject worker manually establishes the classification relationship between an entry and a subject term ("entry" and "equestrian" in the example of fig. 10), the entry contributes to the subject term's definition. For the related icon classes this may mean that the icon class can be derived from or validated by the icons of the entries in an extent;
- vice versa, icons assigned to a description can be matched against an icon class which in return corresponds to a subject term. Using some distance function the SOWing system can derive or propose the classification of an entry.

In this way the subject classification differs substantially from object classification: in contrast to an object class a subject class does not define a uniform structure for all members of its extent. In addition, subject class entries may be members in more than one extent simultaneously and may change subject class membership during their lifetime. In contrast, an entry in its role as object belongs to exactly one object class and this membership is immutable over time.

4.3 Personalization Facilities

We substantiate some of the architectural decisions of the SOWing environment by giving examples from the Warburg Electronic Library. First we look at a the personalization [24] of a description work. Imagine the user downloads a case entry as the following XML document:

```
<picture-card>
  <title>Bonaparte Crossing ...</title>
  <artist>David, Jacques-Louis</artist>
</picture-card>
```

Once a user has copied the description into his personal working environment he is free to modify it at will and might end up with a document like the one shown below.

```
<picture-card>
  <title>Napoleon Crossing ...</title>
  <artist>/artistdb/artist4368.xml</artist>
  <medium>Painting</medium>
</picture-card>
```

Clearly, three changes of a different nature were performed:

- a value change: from "Bonaparte" to "Napoleon";
- a type change: "artist" now is a reference, no longer a value;
- an additional feature is added: "medium".

The latter two changes exemplify the semi-structured nature of the descriptions: part of the personalized descriptions conforms to the schema the community chose for descriptions. Another part was added by the user, freely choosing a tag name, the elements content format, and the position where to integrate the new element. This is the kind of liberty the subject workers expect to have in the SOWing environment.

When the community decides to accept and re-integrate the user's contribution into the community's subject index it has to perform the data (eventually also the schema) update [4]. The SOWing server discovers such type changes by tracking the description work [11].

4.4 Case Entry Generation from XML Documents

A SOWing user may submit a work case document in the form of the following XML document:

```
<report>
  In the <medium>Painting</medium>
  <title>Napoleon Crossing ...</title> the artist
  <artist>David, Jacques-Louis</artist> depicts
  Napoleon riding ...
</report>
```

Then the SOWing system will start analyzing the above work document and generate the following initial work case entry:

```
<picture-card>
  <!-- according to the original conceptual model -->
  <title>Napoleon Crossing ...</title>
  <artist>David, Jacques-Louis</artist>
  <!-- additional markup recognized -->
  <medium>Painting</medium>
</picture-card>
```

This automatically generated version of a work case entry serves as the basis for the description worker's SOWing task.

4.5 SOWing Interfaces and External Tool Support

SOWing entities – work cases, documents, case entries, and subjects - may vary widely by their degree of formalization ranging from rather informal entities (to be mediated to humans) to fully formalized structures (to be read and processed by machines). Furthermore, SOWing entities may be consumed and produced by external tools thus requiring a general interfacing technology to the world outside the SOWing platform. A more detailed discussion in the context of XML can be found in section five.

Our vision of a SOWing system is to *support* a community of users in a long-term process of sharing, evolving and partially formalizing their understanding of a common application domain – and not to *force* them to use a SOWing system. If such a system is to be used it must be possible to create a wide variety of external representations of the entities and relationships maintained by the system.

This is required for two reasons:

- Users want to import the content from such a system into their own working environment (word processors, multimedia authoring tools, expert systems etc.). For this purpose a portable external representation is needed which can be understood by different software systems – at least in part.
- The second reason for a portable and even human readable format is the necessity to facilitate the exchange of the data and the knowledge behind it between different people. Especially those who do not possess an identical or compatible SOWing system will need such an external representation. We consider the extensible markup language as being useful in both scenarios.

On the other hand the SOWing system itself must also be able to process data from various sources. Therefore, either converters between external data and the internal SOWing data model have to be provided or well-defined interfaces to external representations must be available.

5 XML for Subject-Oriented Work

In the previous sections we developed a scenario for a software environment which supports and records subject-oriented work. As soon as the content maintained by such a system is communicated to others or used cooperatively with different systems, the question arises, how to represent content and knowledge about it in a system-independent manner. The Extensible Markup Language and its standards [2], [15] are intended to enhance content structure and coherence and, by doing so, to improve content production and interchange.

Parallel to our work on the SOWing environment much work has been invested in XML-based standards for various content-related services [1]. However, most of these standardization efforts lack a common application scenario which could prove more than the usefulness of isolated standards, e.g., by making two or more of them collaborate. We regard our SOWing system as a cooperative platform for suites of XML-based services.

5.1 XML and Documents

One application of XML in the scenarios mentioned above is its use as a platform-independent notation for content exchange. XML can be read and processed by persons as well as by machines. *Processing* in this context means that machines are able to store XML documents and increase their coherence by performing certain consistency checks, e.g., whether XML data are well-formed or valid according to a certain document type definition. With pure XML, as stated in [10], only syntactical consistency and interoperability can be achieved. If a target machine is supposed to go beyond that point reason about the content of an XML document, domain knowledge has to be hardwired into the processing machine and this knowledge has to be synchronized with assumptions from the content author.

Under this assumption pure XML then could be used to facilitate communication between partners who share a common understanding of the nature of the content being exchanged.

5.2 XML and Semantics

Up to a certain degree XML can be used to represent subject terms and their relationships. However, as [12] argues, important ontological relationships (e.g., "subclass-of" or "instance-of") can not be modeled directly by a XML document type definition. Another problem arises because XML allows different ways to model the same relationship and no support for some notion of equivalence. Properties of a concept can be modeled in at least two ways: as an XML element on its own or as an attribute of another element. Supposing a document type description is used for some ontological statement, the receiver of a document based on this statement has to "know" that this document type represents an ontology and how it can be derived from the document type definition. Currently several proposals for a common formalism for ontology representation are being discussed [10], [12].

The Resource Description Framework (RDF, see, for example, [15]) provides primitives which facilitate the representation of ontologies in a much more natural way compared to (pure) XML. RDF can be implemented on top of XML and is suggested to be used as a basic framework for the definition of common ontology interchange languages, for example, OIL [7].

We expect that by working within the SOWing environment in combination with the above standards we can substantially improve the re-use of publication work and thus simplify the overall publication process. Furthermore, the understanding of documents for readers outside the community for which a work was originally published can be improved and machine reasoning about the works becomes feasible to the extent to which as both, the SOWing system and the processing machine, share subject index information.

5.3 XML and Activities

As mentioned earlier our notion of "work" refers not only to the result of a product process, i.e., to documents on certain concepts or ideas, but to the entire process by which the document is created. Through its reflective capabilities to observe the production process the SOWing system can contribute to questions such as

- in which sequence were the production steps carried out ?
- how long did each production step take ?
- what are the influences leading to a concrete object production sequence ?

For works belonging to the past this information usually is not available or has to be derived from other sources. For description work as carried out by the Art Historians in the WEL project, however, information about the production process of such a description work can be collected almost automatically by the SOWing environment. Work cases being "first class citizens" can themselves become subjects to later research and one does not need to "guess" what the circumstances of a production may have been but one *knows* what these circumstances were – at least to the extent the SOWing environments model documents them.

In this sense the SOWing environment can document the process which led to a certain document and enhance the understanding of the "work" behind the document. On the other hand this process description can be used to derive recommendations for future production processes in the sense of instructions that have to be carried out for tasks similar to a successfully completed process [14]. Well-documented production processes can serve as templates for future production and contribute to process standardization [25]. XML-based languages such as XRL (eXchangeable Routing Language [31]) can be used to represent such processes.

In summary, we expect a twofold contribution when using XML as a common SOWing interface: XML-based tools will substantially and rapidly enhance the SOWing functionality and, in reverse, the SOWing model will give a semantic underpinning and connectivity to XML services and thus significantly improve their usability.

6 Summary and Outlook

Our SOWing platform, experiments and the SOWing project as a whole aim at relating, organizing and defining subjects and documents as well as the content-oriented work behind it. In an interdisciplinary project with our Art History colleagues who are working in the subject area of "Political Iconography" we gained substantial insight into their *Subject-Oriented Working* (SOWing) needs and into initial requirements for a generic SOWing platform. In this paper we outlined the project, its basic models, their generalization as well as our initial experiences with prototypical SOWing implementations and compared our work with various XML-related activities.

On the modeling level we improved our understanding of

- the basic SOWing entities and their relationships;
- the notion of work, i.e., the production context of content;
- the role of key icons and keywords for content-based subject definition.

On the system level the SOWing project is currently investigating the requirements for

- a generator-based architecture for SOWing entities and relationships;
- reflective system technology and its use for advanced SOWing services;
- customized and personalized SOWing indices;
- XML-based tool interoperability.

Future large scale content-oriented project work will have to interact with a substantial number of SOWing indices and, therefore, requires a technology for "plugable SOWing arrays". SOWing indices have to deliver their content through a wide variety of document types ranging from media documents laid out for human interaction to structured (and typed) documents for machine consumption. Finally, SOWing support for the modeling, management and enactment of content-oriented work has to be further improved.

Acknowledgement

We would like to thank the Ministry for Science and Research, County of Hamburg, and the Warburg Stiftung, Hamburg, for their continuous support of our research. In addition we gratefully acknowledge a research grant from the Deutsche Forschungsgemeinschaft, DFG, in support of our R&D project on "Cooperative Reference Libraries".

References

1. Berners-Lee, Tim, Fischetti, Mark: Weaving the Web; the original design and ultimate destiny of the World Wide Web by its inventor. Harper, San Francisco (2000)
2. Bray, Tim, Paoli, Jean, Sperberg-McQueen, C. M., Maler, Eve: Extensible Markup Language (XML) 1.0 (2nd Edition). http://www.w3.org/TR/2000/REC-xml-20001006 (2000)
3. Bruhn, M.: The Warburg Electronic Library in Hamburg: A Digital Index of Political Iconography. In: Visual Resources, Vol XV (1999) 405-423
4. Buckingham Shum, Simon: Negotiating the Construction and Reconstruction of Organizational Memories. Journal of Universal Computer Science, vol. 3, no. 8 (1997) 899-928
5. CoreMedia AG: Homepage. http://www.coremedia.com (2001)
6. Coulter, Neal, French, James, Glinert, Ephraim, Horton, Thomas, Mead, Nancy, Rada, Roy, Ralston, Craig, Rodkin, Anthony, Rous, Bernard, Tucker, Allen, Wegner, Peter, Weiss, Eric, Wierzbicki, Carol: ACM Computing Classification System 1998: Current Status and Future Maintenance. Technical report, http://www.acm.org/class/1998/ccsup.pdf (1998)
7. Cover, R.: Ontology Interchange Language. http://xml.coverpages.org/oil.html (2001)
8. De Michelis, Giorgio, Dubois, Eric, Jarke, Matthias, Matthes, Florian, Mylopoulos, John, Schmidt, Joachim W., Woo, Carson, Yu, Eric: A Three-Faceted View of Information Systems. In: Communications of the ACM, 41(12) (1998) 64-70
9. Deacon, Terrence W.: The Symbolic Species. The co-evolution of language and the brain. W. W. Norton & Company, New York/London (1997)
10. Decker, S., van Harmelen, F., Broekstra, J., Erdmann, M., Fensel, D., Horrocks, I., Klein, M., Melnik, S.: The semantic web: The roles of xml and rdf. In: IEEE Expert, 15(3) (2000)
11. Dourish, P., Bellotti, V.: Awareness and Coordination in Shared Workspaces. In: Proceedings of ACM CSCW'92 Conference on Computer-supported Cooperative Work, ACM-Press (1992) 107-114
12. Fensel, D.: Relating Ontology Languages and Web Standards. In: Informatik und Wirtschaftsinformatik. Modellierung 2000, Foelbach Verlag (2000)
13. Gruber, T. R.: A translation approach to portable ontology specifications. Technical Report KSL 92-71, Computer Science Department, Stanford University, California (1993)
14. Khoshafian, Setrag, Buckiewicz, Marek: Introduction to Groupware, Workflow, and Workgroup Computing. John Wiley & Sons, Inc., New York (1995)
15. Lassila, O., Swick, R. R.: Resource Description Framework (RDF) Model and Syntax Specification. Recommendation. W3C. http://www.w3.org/TR/REC-rdf-syntax/ (2000)
16. Matthes, F., Schröder, G., Schmidt, J.W.: Tycoon: A Scalable and Interoperable Persistent System Environment. In: Atkinson, Malcom P., Welland, Ray (eds.): Fully Integrated Data Environments, ESPRIT Basic Research Series, Springer-Verlag (2000) 365-381
17. Maurer, Hermann, Lennon, Jennifer: Digital Libraries as Learning and Teaching Support. In: Journal of Universal Computer Science, vol. 1, no. 11 (1995) 719-727

18. McLuhan, Marshall: Understanding Media. The Extensions of Man. The MIT Press, Cambridge/Massachusetts/London (1964, 1994)
19. Niederée, C., Hattendorf, C., Müßig, S. (with J.W. Schmidt und M. Warnke): Warburg Electronic Library – Eine digitale Bibliothek für die Politische Ikonographie In: uni-hh Forschung, Beiträge aus der Universität Hamburg, XXXI (1997) 6-16.
20. Nürnberg, Peter J., Schneider, Erich R., Leggett, John J.: Designing Digital Libraries for the Hyperliterate Age. In: Journal of Universal Computer Science, 2 (9) (1996) 610-622
21. Raulf, M., Müller, R., Matthes, F., Scheunert, K.J., Schmidt, J.W.: Subject-oriented Document Administration for Internet-based Project Management (in German). In: Proceedings "Management and Controlling of IT-Projects", dpunkt.verlag, Heidelberg (2001)
22. Rostek, Lothar, Möhr, Wiebke, Fischer, Dietrich: Weaving a Web: The Structure and Creation of an Object Network Representing an Electronic Reference Work. In: Fankhauser, P., Ockenfeld, M. (eds.): Integrated Publication and Information Systems. 10 Years of Research and Development at GMD-IPSI, Sankt Augustin: GMD (1993) 189-199
23. Schmidt, J. W., Matthes, F.: The DBPL Project: Advances in Modular Database Programming. Information Systems, 19(2) (1994) 121-140
24. Schmidt, J.W., Schröder, G., Niederée, C., Matthes, F.: Linguistic and Architectural Requirements for Personalized Digital Libraries. In: International Journal on Digital Libraries, 1(1) (1997)
25. Schmidt, Joachim W., Sehring, Hans-Werner: Dockets: A Model for Adding Value to Content. In: Akoka, Jacky, Bouzeghoub, Mokrane, Comyn-Wattiau, Isabelle, Métais, Elisabeth (eds): Proceedings of the 18th International Conference on Conceptual Modeling, volume 1728 of Lecture Notes in Computer Science, Springer-Verlag (1999) 248-262
26. Schmidt, Joachim W., Sehring, Hans-Werner, Warnke, Martin: The Index for Political Iconography and the Warburg Electronic Library (in German). In: Pompe, Hedwig, Scholz, Leander, "Archiving Processes", "Mediology" Series, Dumont (2001)
27. Simone, Carla, Divitini, Monica: Ariadne: Supporting Coordination through a Flexible Use of the Knowledge on Work Processes. In: Journal of Universal Computer Science, vol. 3, no. 8 (1997) 865-898
28. Smith, John Miles, Smith, Diane C. P.: Database Abstractions: Aggregation and Generalization. In: TODS 2(2) (1977) 105-133
29. Sowa, John F.: Knowledge Representation, Logical, Philosophical, and Computational Foundations. Brooks/Cole, Thomson Learning (2000)
30. Svenonius, Elaine: The Intellectual Foundation of Information Organization. The MIT Press, Cambridge/Massachusetts/London, England (2000)
31. van der Aalst, W. M. P., Kumar, A.: XML Based Schema Definition for Support of Interorganizational Workflow. http://tmitwww.tm.tue.nl/staff/wvdaalst/Workflow/xrl/isr01-5.pdf
32. van Waal, Henri: ICONCLASS – An iconographic classification system. North-Holland Publishing Company, Amsterdam/Oxford/New York, completed and edited by L.D. Couprie, R.H. Fuchs, E. Tholen, G. Vellekoop, a.o. (1973-85)
33. Warnke, Martin: Bildindex zur politischen Ikonographie. Forschungsstelle Politische Ikonographie, Kunstgeschichtliches Seminar der Universität Hamburg (1996)
34. Warnke, Martin (ed.): Der Bilderatlas Mnemosyne. Unter Mitarbeit von Claudia Brink. Akademie Verlag, Berlin (2000)
35. Webster's Third New International Dictionary of the English Language, Chicago (1996)