
From Information-Centric to Experiential Environments

Rahul Singh¹ and Ramesh Jain²

¹ San Francisco State University, San Francisco, USA

² University of California, Irvine, USA

Summary

With progress in technology, information management systems are transitioning from storing well defined entities and relationships to the challenge of managing multifarious heterogeneous data. Underlying such data is often a rich diversity of information with emergent semantics. Recognizing this characteristic is essential to executing the transition from data to knowledge. In this context, this chapter presents the paradigm of experiential environments for facilitating user-data interactions in information management systems. Experiential environments emphasize obtaining information and insights rather than pure data lookup. To facilitate this aim, the paradigm utilizes the sentient nature of human beings, their sensory abilities, and interactive query-exploration-presentation interfaces to experience and assimilate information.

1. Introduction

In the good old days, just a decade or so ago, to exemplify the requirements and structure of a database, we typically considered a corporate database. Within it, entities, such as “employee”, “address”, or “salary” consisted of alphanumeric fields. Each such field represented some attribute that had been modeled. Users would pose queries, for example, to discover an employee attribute or to find all employees that satisfied certain attribute-related predicates.

Although in the new millennium users have vastly different expectations, most databases still retain the design philosophy of yesteryears: *Users ask queries to get answers in an information-centric environment*. This premise holds as long as all users have same or similar requirements. The

Rahul Singh and Ramesh Jain

database can then act as a resource that provides a well-defined environment for articulating queries on a fixed information structure. However, the emergence of Internet based systems, including the WWW, and progress in computing and storage technologies has fundamentally changed, the kind of data that is in common use today. This change is both quantitative and qualitative and has important consequences for paradigms used for interacting with data. Taken together, the situation contrasts sharply with the scenario that was common even just a decade ago and creates a mismatch between existing design paradigms and evolving requirements. A simple analysis of the nature of the data and expectations of users from current and next generation systems highlights the emerging issues:

- Volume of data is growing by orders of magnitude every year.
- Multimedia and sensor data is becoming more and more common.
- Different data sources provide facets of information which have to be combined to form a holistic picture.
- The goal of data assimilation increasingly requires that spatio-temporal characteristics of the data are taken into account
- In many applications, real-time data processing is essential
- Exploration, not querying, is the predominant mode of interaction, which makes context and state critical.
- The user is interested in experience and information, independent of the medium and the source.

In the chapter, we explore the paradigm of experiential computing for designing information management systems. The idea of experiential computing is built on the fact that humans are sentient observers. Therefore, this paradigm emphasizes interactivity and support for experiential user factors in the quest for information assimilation.

1.1 Issues Motivating the Need for Experiential Computing

Let's look at some of the basic issues that underlie this change. We motivate our perspective by noting three critical factors that influence the situation. These include: How the data is modeled and its implications (the data model), the presence of different data types and their implications (data heterogeneity), and finally, the expectation of users as they interact with the data (user requirements).

Data model: A data model can be thought of as an abstraction device through which a reasonable interpretation of the data can be obtained [53]. Keeping this in mind, we can identify two types of data sources, those that

From Information-Centric to Experiential Environments

are *strongly modeled* and those that are *weakly modeled*. Conventional databases such as an inventory database or a corporate database are examples of strongly modeled sources of data. In them, data and relationships amongst data are stored with very specific goals in mind. Information from such databases can be retrieved (in the sense of normal queries) only in terms of the data and relations that are explicitly modeled or derivations based on them (e.g. joins on tables). Weakly modeled data sources on the other hand are less specific about the interpretations of the data that are made available through them. A general web-page is an example of a weakly modeled data source. While thematically coherent, a general web-page, unlike the above examples, does not seek to limit access to information present in it through a limited set of entities and relations.

Data heterogeneity. Data in traditional databases is synonymous with alphanumeric information. However, many contemporary applications are characterized by the fact that information is represented through different types of data like text (alphanumeric), images, audio, video, or other specialized data types. The proliferation of physically different data types (or data heterogeneity) is driven both by the increasing capabilities of computational systems as well as the increasing sophistication and ease of use of digital sensor technologies coupled with their decreasing cost. One of the important challenges in situations that involve heterogeneous data types lies in that the semantics associated with complex media (like video or images) is *emergent*, i.e. media is endowed with meaning by placing it in context of other similar media and through user interactions [44]. This has strong bearing on the systems that are designed to work with such data. For instance, the emergent semantics of complex media imply that such information is necessarily weakly modeled. To capture such issues, we distinguish two types of data: *alphanumeric* and *media-based*, where the latter may include alphanumeric data when occurring in conjunction with other data types.

User requirements: User requirements for the data fall into roughly two categories: *information* and *insight*. For example, in some cases, a user is just looking for some information, such as the location of a specific restaurant. In other cases the user may be interested in more complex insights such as how cosmopolitan is a particular city. These two types of requirements present completely different set of challenges for the design of information management systems.

1.2 Towards a New Paradigm

The matrix in Fig. 1 captures the relationships between the aforementioned issues. Each cell of this matrix lists the paradigms which can be used to address requirements at the intersection of these issues.

		Experiential Environments	
Insight		<i>Visualization (Indirect Experiential Environments)</i> <i>Data Mining</i>	<i>Direct (heterogeneous data) or Indirect (alphanumeric data) Experiential Environments</i>
	Information	<i>Current Databases</i>	<i>Search Engines</i>
		Strongly Modeled Sources	Weakly Modeled Sources

Fig. 1. Paradigms at the intersection of data modeling, data heterogeneity, and user requirements

In this matrix, predictably, databases lie in the lower left quadrant at the intersection of information and strongly modeled sources; they are ideally suited for obtaining precise information in well specified domains. The bottom right quadrant is occupied by search engines. They are well tailored for generic searches across weakly modeled information sources. Such sources may either be alphanumeric or have heterogeneous data typically with textual annotations. The primary intention of search engines is to provide information through responses to specific (keyword-based) queries and not to directly facilitate exploration.

In this matrix, predictably, databases lie in the lower left quadrant at the intersection of information and strongly modeled sources; they are ideally suited for obtaining precise information in well specified domains. The bottom right quadrant is occupied by search engines. They are well tailored for generic searches across weakly modeled information sources. Such sources may either be alphanumeric or have heterogeneous data typically with textual annotations. The primary intention of search engines is

From Information-Centric to Experiential Environments

to provide information through responses to specific (keyword-based) queries and not to directly facilitate exploration.

The top half of Fig. 1 consists of paradigms, many of which demonstrate characteristics that are partially or wholly experiential. This transition is reflected in the top left quadrant which consists of approaches that support gaining insights from precise sources. Techniques commonly employed for this goal include data mining and visualization. The latter is of special interest to us as it seeks to transform and present data in a manner that allow users to gain insights by “seeing” the patterns and relationships that may be present. This attribute, where the human senses are involved directly, as the reader will see from the ensuing sections, is a key characteristic of the experiential paradigm.

Finally, the top right quadrant addresses the intersection of insight and imprecise data sources. This intersection produces challenges which can be addressed through *experiential environments*, a new way of interacting with data that will become increasingly common in most data-intensive applications: In such cases, users encounter immense volumes of multifarious data from disparate, distributed, sometimes even unknown, data sources. To gain insights from such data, one must be immersed in it, just as one would be immersed in a real life situation to experience it first hand. Humans are sentient observers. They want to explore and experience information. Furthermore, they typically prefer to directly interact with the data without complicated intervening metaphors. This tendency probably stems from the fact that we – humans are all immersed in the real world – where the real world is really different attributes at different points surrounding us. We use our senses to measure or infer the various attributes. For example, our visual system is a powerful mechanism that allows us to infer different kinds of attributes about the environment surrounding us. Similarly, tactile senses allow us to measure other characteristics of the environment that are in close proximity to us. As these examples illustrate, we have complementary sensors to facilitate our explorations and experiences. Vision and sound are our sensors to infer about the world without the constraint of close proximity and we use these for experience as well as for communications. Other sensors, like touch, are used in situations where a certain amount of proximity and intimacy is required.

We develop the ideas initially proposed in [28] to point out two types of experiential environments that may be contemplated: The first of these is the *indirect experiential environments*. Within these environments, data is transformed to present it in manners where users can involve their senses to discern patterns and relationships. Techniques in information visualization as well as more evolved and integrated approaches such as the business activity monitoring application covered in Section 4.2 of this chapter

fall in this category. The second type of experiential environment is called a *direct experiential environment*. The fundamental difference of such environments from those in the first category lies in their ability to deal with data types such as imagery or audio that can be directly presented to users. Therefore there is no interpretation or selection of transformations involved (Section 4 describes an example of such a system directed at the problem of personal multimedia information management). Finally, user-information interactions are also direct in that they do not use any intermediate metaphors or transformations.

1.3 From Data to Information and now to Insights: The Etudes of Experiential Computing

There is a very clear trend in the evolution of computing approaches from databases to search engines. Belew [2] presents this trend by comparing the key characteristics of these systems. In Table 1, we extend Belew's observations to include experiential environments. Our extension highlights the trend from data to information and now to insights.

Table 1. Data to Information (and now) to insights

	Database	Information Retrieval	Experiential Environment
Basic goal	Provide data	Provide information sources	Provide Insight
Data Type	Alphanumeric	Text	Multimedia
User Query	Specific	General	Emergent
System Provides	Data item	Pointer	Heterogeneous data
Retrieval Method	Deterministic	Probabilistic	Hybrid
Success Criterion	Efficiency	Utility	User Satisfaction

Traditional databases were designed to provide an environment where user could articulate their information needs using precisely specified logi-

From Information-Centric to Experiential Environments

cal relationships. The database would then respond by providing the information. On-Line-Analytical-Processing (OLAP) and visualization-based approaches are based on the same systems, but go farther by pulling out a volume of data and then using visualization tools to allow exploration of the retrieved dataset. Search engines directly adopted the basic concept of query from databases. Thus, in most of the current systems, a user articulates a query and gets an answer for it. If further information is needed, a new query must be articulated. Current information environments actually work against the human-machine synergy. Humans are very efficient in conceptual and perceptual analysis and relatively weak in mathematical and logical analysis; computers are exactly the opposite. In an experiential environment, users *directly use their senses to observe data and information of interest related to an event and they interact naturally with the data based on their particular set of interests in the context of that event*. Experiential environments have the following important characteristics:

- *They are direct*: These environments provide a holistic picture of the event without using any unfamiliar metaphors and commands. Within them, users employ intuitive actions based on commonly used operations and their anticipated results. In experiential environments, a user is presented data that is easily and rapidly interpreted by human senses and then the user interacts with the dataset to get a modified dataset.
- *They provide the same query and presentation spaces*: Most current information systems use different query and presentation spaces. Consider popular search engines. They provide a box to enter keywords and the system responds with a list of thousands of entries spanning over hundreds of pages. A user has no idea how the entries on the first page may be related to the entries on the last, or how many times the same entry appears, or even how the entries on the same page are related to each other. Contrast this to a spreadsheet. User articulates a query by changing certain data that is displayed in context of other data items. User's action results in a new sheet showing new relationships. Here the query and presentation spaces are the same. These systems are called What-You-See-Is-What-You-Get or WYSIWYG.
- *They consider both the user state and context*: Information system should know the state and context of the user and present information that is relevant to this particular user in the given state and context. People operate best in known contexts and do not like instantaneous context switching. Early information systems, including databases, were designed to provide scalability and efficiency. These considerations led

Rahul Singh and Ramesh Jain

to designs that were state-less. The efficiency of relational databases is the result of this decision. This is also the reason why Internet search engines, which do not store user states, can be dissatisfying, as users seek to drill-down on information obtained from previous queries.

- *They promote perceptual analysis and exploration:* Experiential systems promote perceptual analysis and exploration. Because users involve their senses in analyzing, exploring, and interacting with the system, these systems are more compelling and understandable. Text based systems provide abstract information in visual form. Video games and many simulation systems are so engaging because they provide powerful visual environment, sound, and in some cases tactile inputs to users.

In this chapter, we begin by discussing the data engineering challenges that underlie the development of experiential environments in Section 2. This is followed by a description of event-based modeling in Section 3. In this section, we discuss how event-based organization and management of data facilitates development of contextual and personalized experiential environments. In Section 4, we present descriptions of two experiential systems in the areas of personal information management and business activity monitoring respectively. These examples illustrate how the ideas espoused in the previous sections can be realized in designing real-world systems. We conclude this chapter in Section 5 by reiterating the fundamental ideas behind the paradigm of experiential environments and outlining its broad applicability in the evolution of the next generation information and data management systems.

2. Data Engineering Challenges for Designing Experiential Environments

Experiential computing environments require supporting user interactions such as browsing, exploration, and queries on information represented through different media. The *direct* nature of experiential environments implies that the results of such interactions are expressed in the native format(s) of the underlying data. The data engineering challenges that are encountered in designing such systems span issues related to modeling and representation of heterogeneous data as well as design of user-data interfaces that support interactions that are direct and user context aware.

In experiential computing, the problem of *heterogeneity*, arises in many forms, including: *infrastructural heterogeneity* (due to different types of hardware and software platforms that may be involved), *logical heteroge-*

From Information-Centric to Experiential Environments

neity (arising out of different data models or schemata used for providing a logical structure to the data), and *physical heterogeneity* (owing to the presence of fundamentally distinct types of data such as text, audio, images, or video). Amongst these, the problem of logical heterogeneity has typically been considered in traditional database research, while that of physical heterogeneity has been the focus in multimedia data modeling.

It can be postulated that physical heterogeneity leads to logical heterogeneity, since the need to capture specificities of each media would result in the development of different data models and schemata, which ultimately need to be integrated. This is reflected in the similarities that can be discerned between the approaches to addressing heterogeneity in database and multimedia research. For example, the idea of using local-as-view and global-as-view approaches [25] for specifying the correspondences between data at the source and in the global schema bears parallels to the principles of autonomy and uniformity suggested in [51] for media integration. However, till date, the issue of integration, when the heterogeneity is due to different media, has typically remained unaddressed both in multimedia and database research, as techniques have tended to concentrate on issues that arise in media-specific management. The result of this research focus manifests itself today in a large number of media specific solutions, such as those for images, audio, or video, but hardly any, that span different media.

Our emphasis in this chapter is on the problem of dealing with heterogeneity by starting from the physical level. Consequently, the following sections explicitly deal with issues that arise when information is represented using multiple and distinct media. As explained above, such a formulation subsumes the problem of dealing with logical heterogeneity. Furthermore, this approach allows us to address issues that occur due to the increasing availability of sensor-based data in applications varying from personal information management, to ubiquitous computing and sensor networks. We refer readers who are specifically interested in the issues of managing logical heterogeneity to introductory material well codified in textbooks such as [15, 38] along with the reviews [27, 31] and references therein. In this context, we also emphasize that the principles of Experiential Computing are not limited to the availability of multimedia information and are equally applicable to domains where the data is alphanumeric. In Section 4 this is illustrated through two examples, one of which deals with multimedia data in the domain of Personal Information Management and the other, with alphanumeric data, in Business Intelligence.

2.1 Understanding (Physical) Heterogeneity

Fundamentally, multimedia data has a *gestalt* nature. This implies that it is comprised of more than one media that are semantically correlated and that the complete semantics conveyed in the data can not be discerned by considering the data streams individually. The classical example of capturing an explosion using image-based data (video) and sound (audio) is often forwarded to underline this aspect. In the case of this example, either of the media (a flash of light or a loud sound) taken in isolation, is insufficient to determine that an explosion occurred. A unified data model is thus essential not only for preserving the semantic integrity of the information, but also for conveying the same by supporting query-retrieval and user interactions with the data. The principal characteristics of multimedia data that have influence on its modeling and subsequent usages include:

- *Semantic correlations across media:* As briefly described above, representing the semantic correlations across different media is fundamental to storage, processing, query-retrieval, and utilization of multimedia information. How do we represent heterogeneous multimedia data in a *general* and *unified* way that emphasizes the semantic correlations between the media? Different media, such as audio and images, have different form and characteristics. A unified data model needs to address these issues. This issue also expresses itself in terms of the problem of *multiple representations* where a single object, entity, or phenomenon may be captured and represented in different media formats. For example, the state of a patient's health may be recorded using bio-medical imagery data including time-varying imaging, alpha-numeric data detailing blood pressure, body-temperature, or patient-weight, and audio-transcripts. A data model should be capable of seamlessly resolving across the various representations.
- *Temporal characteristics of the data:* From surveillance videos, to personal photographs, to biomedical imagery (for example, tracking of synaptic activation in the brain), temporal and dynamic phenomena are commonly represented through multimedia data. The aforementioned applications and many others are characterized by data in which the time of occurrence or changes over time denote valuable information. A multimedia data model, therefore, should be able to represent, query, and reason about time and time-based relationships. Directly using traditional data models, such as the *entity-relationship* paradigm, to reason about dynamic data is complicated because such an approach is primarily designed to reflect a set of static relationships between entities. In dynamic environments not only the attributes associated with entities,

but also the relationships between entities change with time. In modeling multimedia data, an additional challenge is faced in situations that require integrating semantically related dynamic and static media. This can happen for example, when information about a sporting event is available through a video-recording and a text-report.

- *Spatial characteristics of the data:* Much of multimedia data has inherent characteristics that can be correlated with location. For example, automated traffic monitoring at intersections provides video footage distinct from those taken at highways. In bio-medical imaging the location (organ or tissue) is of critical importance for analysis and interpretation of the images. Personal videos of trips can be categorized by the geographical locations visited. These examples indicate that the semantics, the form, and the relationships expressed in multimedia data are often influenced by location and relationships such as adjacency, connectivity, proximity, or containment that can be defined over space.
- *User interactions with the data:* Current multimedia systems typically consider information in a manner that is independent of the user and context. Further, they make an implicit assumption that acquisition of knowledge by the user (based on the media) is a linear process and can be adequately represented by the rendering of the media alone. Therefore, they support limited interactions between the user and the data. This causes a significant loss in the totality of information being communicated between creation of the media and its consumption by users, as the media is presented in isolation from the context of its creation [49] as well as the context of the user. For instance, the rendering of a video is typically from a single perspective and users have little or no control in interactively exploring a scene using multiple perspectives. Additionally, in true *multi-media* settings, disparate data sources need to be united for presentation, query, and exploration in manners where the users are free to state their requests in natural form based on objects and event relationships of interest. The entire set of aforementioned issues is complicated by the fact that the semantics associated with complex media (like video or images) is emergent. Developing user-data interaction paradigms that address such issues requires support from underlying data models to impart the appropriate structure to the information.

2.2 Previous Research in Multimedia and Databases

In recent years, a number of data models have been developed to address the structure and semantics of media data like images and video [9, 10, 11, 12, 16, 20, 21, 33, 37, 46, 50] or sound [7, 55]. Such research has typically

focused on the development of powerful features to describe the corresponding media and the use of similarity functions to answer queries based on these features [43]. This approach simplifies the general multimedia database problem, because a database is assumed to contain only a specific type of media data [13]. A related line of research has focused on developing models that support the structure of media data and the syntactic operations that are typically performed on them. For example [19] proposes as a basic abstraction the notion of “timed streams” of media elements like video, audio samples, and musical notes. This model considers issues like the temporal nature of media data (defined in terms of real time presentation constraints and media synchronization) along with operations like media derivation and media composition. Similarly [34] considers temporal access control issues like reverse, fast-forward, and midpoint suspension/resumption in their model. An object-relational model that builds upon [19] is proposed in [13] where a three-layered structure is defined. The lowermost layer consists of raw data (byte sequences). The middle layer consists of multimedia entities, called multimedia type, which can be images, image stacks, sound, video, or text. The top layer consists of logical entities that model the domain semantics and interact with entities representing multimedia types. Additionally, specialized entities, containing visual and spatial information can be defined at the top layer to be used for content-based querying. Layered architectures such as [56] or the one described above, break down the complexity of multimedia modeling by seeking clear distinction between raw data modeling, conceptual modeling, and presentation management.

The *Garlic* project at IBM [8, 24], uses an object-oriented data model as middleware to integrate multiple (potentially multimedia) databases. Translations of data types and schemas between individual repositories and Garlic are accomplished using repository wrappers. A repository of complex objects is provided for integration of multimedia data and legacy data. Query processing and data manipulation are supported through the Garlic system.

The papers reviewed above catalog a rich diversity of research approaches towards modeling information represented through multimedia data. However, to the best of our knowledge, *no research till date has attempted to address within a single framework, the problem of multimedia information management, in context of all the issues we enumerated early on in this section.* Towards this, in the following section we introduce the rudiments of event-based modeling that form the basis of our research in designing experiential systems.

3. Event-Based Unified Modeling of Multimedia Data

3.1 The Conceptual Model

The fundamental idea underlying the data model being considered by us is the notion of an *event* which may be defined as under [47]:

Definition 1: An event is an observed physical reality parameterized by space and time. The observations describing the event are defined by the nature or physics of the observable, the observation model, and the observer.

Certain key issues in this definition need to be highlighted. First, events are treated as a fundamental physical reality and the observations (or data) that describe them are defined to depend on the observation model and the observer. The observation model may include among others, the observation method (e.g. audio, video, images, or other data types such as alpha-numeric data), sampling model (e.g. video-rate), and sampling period. The role of the observer is fulfilled by users involved in creation or consumption of the information. Events thus constitute the unifying notion that brings together heterogeneous data that is semantically correlated. This idea is illustrated in Fig. 2. Using events as the central semantic notion, a conceptual model can therefore be developed. As part of the conceptual model, the specification of an event covers three primary aspects:

- *Event information:* The information component of the event may consist of specific attributes. Since events are spatio-temporal constructs, the event information component necessarily contains the time period of the activity and its spatial characteristics, e.g. its location. Additionally, information required to uniquely identify an event are also stored here. Further, entities like people or objects that participate in an event may be described here along with other types of domain specific information.
- *Event relations:* Events (and the activities underlying them) may be related to other events (activities) that occur in the system. Examples of such relations can be temporal and spatial co-occurrences, temporal sequencing, cause-effect relations, and aggregations of events. This information is modeled and described in the event relations component.
- *Media Support:* Each event is a unifying point for the various observations that describe it. These observations are available to us through different types of media data. Specific media data is said to support an event, if it captures (describes) that event. We note that the exact form of the description depends on the characteristics of the media. For example, a basketball game may be described by video, photographs, and

Rahul Singh and Ramesh Jain

mixed text-image new article. Each of these descriptions exemplifies specific media that have different characteristics, while supporting the same event. Such media data may reside as multimedia files in a file system or in media specific databases. In the media support component, information such as media types, resource locators, or indexes corresponding to specific media that support the given event are stored. It should be noted that the conceptual model imposes no restrictions on the same media simultaneously supporting multiple events.

3.2 Multi-Level Events

For broad usability, in the context of media-rich data, a data model and its physical implementations need to incorporate domain-level semantics so that users can interact with the information at the semantic level rather than at the level of non-intuitive media-specific features. This requires a model to span and seamlessly transition between low-level signal-centric modeling to high-level concept and semantic modeling. As an example of this, consider the problem of mapping an image with large number of “red” colored pixels to an entity with clear semantics associated with it, such as “tomato” or “sunset”. This issue is synonymous with the problem of bridging the signal-to-symbol gap (the gap between the signal-level description of the content and the symbolic meaning associated with it).

One of the primary goals of event-based modeling is to relate media to semantically meaningful entities and concepts. This requires bridging the signal-to-symbol gap. To assist in building models that encompass such a transition, we distinguish between three types of events. We call the first of these *data events*, the second *elemental events*, and the last *semantic events*.

Data events model the physical characteristics of the information. For example, a photograph or video clip consists of pixels that contain illumination-intensity information of the scene. Similarly, an (digital) audio-clip consists of samplings of a sound wave. These are examples of data events. As these examples illustrate, the media support for each data event is a non-empty singleton set, consisting of a specific media instance. The reader will note that given the signal-level nature of data events, issues such as the nature or physics of the observable and the observation model have a strong bearing on their definition of data events.




Fig. 2. Intuition behind the event-based unified multimedia data model

An elemental event on the other hand is a conceptual entity and reflects an interpretation or analysis of information that is partially inspired by domain semantics while retaining the dependence on signal-level information (data events). For example, detecting a “person speaking” based on audio-visual data is an example of an elemental event. While being independent of data characteristics per-se, the formulation of elemental events draws on detection of data features and may form the focal point for multimedia unification. In context of the earlier example, detecting that a person is speaking may be done by analyzing the sound-level (pure audio-based), analysis of lip movements (pure video-based), or through an analysis involving both audio and video. An elemental event therefore, may have either homogeneous or heterogeneous media support based on the constituent data events.

Finally, a semantic event captures the conceptual (or semantic) interpretation of the data and is based on the underlying elemental events. For instance, “giving a speech” is an example of a semantic event. This event may be based on the elemental event “person speaking” and be temporally related with other elemental events such as “coming to the podium”. Thus, semantic events do not have direct media support, but function as a unifying point for all the underlying media.

The reader may note that denotation of data, elemental, and semantic levels in the multi-level event model does not imply that implementations are restricted to have only three levels in the transition from a signal-centric to a semantics-centric modeling of the information. Indeed, in an

implementation, one or more of the levels defined above may consist of sub-levels to assist in the transition. This will especially be true for the highest level where semantic events are represented. Conversely, an implementation may choose to collapse these three levels into a single layer. This could occur, for example, in situations where users directly annotate media to endow high-level semantics to it. In such a case the signal-to-symbol gap is bridged using cognitive input. The system described in Section 4 for personal multimedia information management takes this approach.

3.3 Modeling Time and Space

Time and Space are two of the fundamental attributes of the event model and how they are represented significantly impacts our ability to reason with the proposed model. It may be noted that modeling of time and space has received significant attention in both database and knowledge representation communities (see [54] and references therein) and our approach draws significantly from prior results in the area. In the context of temporal representation, a simple approach is to tag each attribute (or tuple) with a discrete timestamp. Its deficiency lies in that common algebraic operations like addition, multiplication, and division are not applicable to timestamps. Further, information that is not explicitly represented becomes difficult to query.

Research in temporal databases has also explored interval-based models of time. Such representations are ideally suited to describe events (such as a game or a meeting) that occur over a period of time. However, the modeling problem we are considering is significantly more complex and can not be sufficiently addressed through interval-based models only. As an illustration, consider the example of parents taking a digital photograph of the “first smile” of their child. The photograph is in itself an event (a data event, in the three-layered hierarchy), that has an infinitesimal duration (manifested using a single timestamp). Further, based on that single photograph, an interval can not be defined for the (semantic) event “first smile”. In such cases either the fundamental nature of the event or lack of domain semantics precludes the use of interval representations. We therefore propose two temporal data types: infinitesimal time points and time intervals to temporally characterize events. In the following, we denote time points with a lowercase letter, potentially with subscripts (e.g. t_1, t_2) and time intervals with upper case letters $T = [t_1, t_2)$. Algebraic operators can be used to convert information among these types. For example, time intervals can be added or subtracted from time points to yield new time points. Further,

From Information-Centric to Experiential Environments

time points can be subtracted to determine time intervals. Three classes of relationships can then be defined to reason about temporal data. These include:

- *Point-Point Relations*: Assuming a complete temporal ordering, two arbitrary time points t_1 and t_2 can be related as: $t_1 < t_2$ (**Before**); $t_1 = t_2$ (**Simultaneous**); $t_1 > t_2$ (**After**), and the negations $t_1 \geq t_2$ (**Not Before**); and $t_1 \leq t_2$ (**Not After**).
- *Point-Interval Relations*: The relations between an arbitrary time point t_1 and an arbitrary time interval $T = [t_a, t_b)$, are: $t_1 < T \Rightarrow t_1 < t_a$ (**Before**); $t_1 \in T \Rightarrow t_a < t_1 < t_b$ (**During**); $t_1 > T \Rightarrow t_1 > t_b$ (**After**), $t_1 \geq t_a$ (**Not Before**); and $t_1 \leq t_b$ (**Not After**).
- *Interval-Interval Relations*: Given two intervals $T = [t_a, t_b)$ and $U = [u_a, u_b)$, the possible relations between them are [1]: $t_b < u_a$ (**Before**); $t_a = u_a$ and $t_b = u_b$ (**Equal**); $t_b = u_a$ (**Meet**); $t_a < u_a$ and $t_b < u_b$ and $u_a < t_b$ (**Overlap**); $t_a > u_a$ and $t_b < u_b$ (**During**); $t_a = u_a$ and $t_b < u_b$ (**Start**); $t_a > u_a$ and $t_b = u_b$ (**Finish**); and the corresponding symmetric relationships (excluding the case for **Equal**).

These relations allow us to deal with relative position of intervals and are necessary to reason about effects that may influence the occurrence of each other (causality) or manifest themselves with delay. A graphical description of these relations is shown in Fig. 3.

The ability to reason about space, analogous to reasoning about time is a key component in a model that seeks to describe data collected in dynamic settings. Multimedia data, like photographs and videos have obvious spatial (geographic location) characterization associated with them. A wide range of examples from application areas can also be observed. For instance, in biomedical imaging, the location of a tumor in the body is a critical piece of information for diagnostics and treatment. Similarly, in weather forecasting, the localization of severe weather phenomena is critical.

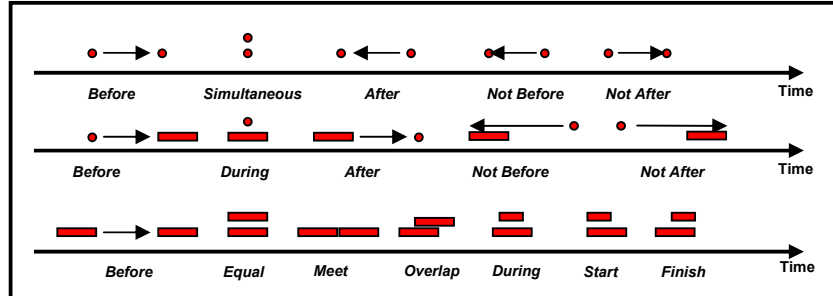


Fig. 3. Illustration of the three types of temporal relationships in the event-based model: The point-point relations are shown at the top, the point-interval relations are shown in the middle, and the interval-interval relations are shown at the bottom

Research in spatial databases suggests two alternative ways of representing space. The first involves describing space itself, i.e. describing every point in space. The second involves representing distinct entities. However, such a bifurcation introduces the problem of reconciling spatial granularity. This can potentially be resolved [23], by supporting concepts for modeling single objects (represented as points, lines, or regions) and spatially related collection of objects (represented as partitions or networks). One unified representation that can support this is the concept of *realm*, introduced in [22], where a realm is defined as a constrained finite set of points and line segments over a discrete grid, and conceptually represents the complete underlying geometry of one particular application space. The constraints ensure the necessity and sufficiency of the grid points for spatial representation. Abstractions such as points, lines, regions, partitions, and networks can either be described either as elements of a realm or represented on top of such elements. This approach appears to hold promise for modeling spatial attributes, by layering events in a hierarchy such that events in each layer share the same semantics of space. Other representations like TIN (triangulated irregular networks) [45] or constraint databases [30] can also be used to address such problems.

To facilitate reasoning with space, spatial algebras or spatial data types need to be defined, so that they capture the fundamental spatial abstractions, the relationships between them, and possible operations on them. In the spatial database community a variety of such approaches have been suggested (see [22, 23, 45] and references therein). Based on these we define the different types of spatial operations that are needed to reason about spatial characteristics of events to include: intersections of spatial types, topological relationships (*containment, intersection, adjacency, and enclo-*

sure), operators defining numeric spatial attributes like *distance*, *area*, arithmetic operators on spatial types (*addition*, *subtraction*) that could be useful in aggregation relationships, operators returning atomic spatial data types (e.g. intersection of lines returning points, union of regions defining regions), operators defining directional relationships like *above*, *below*, *north_of*, and operators defining numeric relationships like *distance < 5*. Researchers in spatial databases have identified various topological relations that may possibly exist. A simplification of these results was suggested in [23] which proposed the following five basic topological relationships derived from intersection of boundaries and interiors: *touch* (defined over line-line, point-line, point-region, line-region, and region-region), *in*, *cross*, *overlap*, and *disjoint*.

4. Putting it all Together: Experiential Environments for Real World Problems

In this section we consider two application domains where issues related to data heterogeneity, the importance of exploration, and the role of user-data interactions in information assimilation play a significant role. These application domains are: multimedia personal information management, and business activity monitoring. We also discuss the design of two systems based on the principles of experiential computing for these domains.

4.1 Application 1: Multimedia Personal Information Management

With advances in processing, storage, and sensor technologies over the last few decades, increasingly digital media of different types is being used to capture and store information. In the specific context of personal information, this trend has significantly accelerated in the recent past with the introduction of affordable digital cameras, portable audio recorders, and cellular phones capable of supporting, capturing, and storing text such as e-mails or instant messages, images, videos, and sound clips. These devices are setting a trend with people capturing increasing amounts of multimedia information to chronicle their day-to-day activities [4]. The emerging area of personal information management seeks to study challenges associated with management, presentation, and assimilation of such information.

4.2 Specificities of Personal Information Management in Media Rich Environments

The very nature of personal information management, especially in media rich environments, introduces specificities that need to be accounted for in a solution methodology. In this context, some key tendencies and requirements that can be gleaned from studies and prior research in the area include:

- *Support for context*: Personal information management systems have to serve the twin functions of *finding* and *reminding*. Rich contextual cues, such as time, space, thumbnails, or previews have been shown to help in search and presentation of personal information [3, 14, 52]. Additionally, user state and context can also be used as powerful aids, as discussed earlier in this chapter. Further, presentation of personal information is typically done not just in terms of isolated media, but by making it part of a specific personal context. Recent systems have attempted to provide support for this notion through the use of concepts such as landmarks (birthdays, deadlines, news events, holidays) [14] and story-lines [16].
- *Co-location of related information*: The necessity of co-locating related information in a system, regardless of their format, can significantly reduce the cognitive load on users [14] and help them in assimilating the information by providing a holistic picture [28].
- *Query versus exploration*: Short (in terms of word length) and simple (in terms of Boolean operators involved) queries have typically been the norm in personal information management systems [14]. It has also been noted, that users of such systems tend to favor navigation and browsing over the use of powerful (but complex) search capabilities [6]. In [28], a review of media-rich applications including personal information, sports, and situation monitoring supports the importance of exploration over pure syntactic querying in forming insights based on the data.
- *Flexible information organization*: Specific media may simultaneously be part of different conceptual organizations defined by a user on the information space [17, 41, 40, 18]. Models, like directories or tables, that enforce rigid data categorization may constrain the way people like to structure and explore information. Such problems can be ameliorated by supporting flexible information organization.
- *Interactive interactions*: Given the multi-modal nature of the available information and the observed user tendencies to eschew complex queries, highly interactive and iterative query strategies are essential for

supporting fruitful user interactions with such systems. Recent efforts have moved in this direction. For example, in [14] and [48] interactive systems for managing personal information is proposed. In [52] the authors propose query-retrieval of digital images using spatial information and interactive queries. In all these cases, the systems emphasize interactive queries, direct presentation of results, and use of contextual cues such as time, participants [14, 48] and location [48, 52]. Evaluations of such systems indicate their efficacy both in terms of quantitative metrics as well as in terms of user satisfaction [14, 48].

4.3 An Experiential Approach to Managing Personal Information: The eVITAE Project

An analysis of the aforementioned problem specificities demonstrates a close relationship between the challenges that pervade personal multimedia information management and the emphasis areas of the experiential computing paradigm. For instance, issues such as context-support and co-location of related information are intimately tied to characteristics of experiential systems such as media independent information modeling and presentation and/or description of spatio-temporal relationships in the data. Similarly, the preference of users for interactive queries and flexible information organization observed in personal information management settings eminently fit the emphasis on interactivity that is central in experiential environments. The goal of the *eVITAE* (electronic-vitae) project [48] is to research the synergy amongst these issues and develop experiential systems for management of personal multimedia data. The prototype consists of three primary components namely *event entry*, *event storage*, and *event query and exploration environment* each of which are described below.

4.4 Event Entry

The role of event entry is to acquire all necessary information to create the event model. For example, such information may include time, location, participants, or any other domain specific event attribute. As has been pointed out [52], the primary ways of acquiring such information include (1) manual entry, (2) from data or data capture devices such as image headers and GPS enabled cameras [42], (3) from a digital calendar, (4) from surrounding information, and (5) by media analysis and association. Currently, *eVITAE* supports the first two approaches. To do so, batch processing scripts have been written to assimilate the media into the database. These scripts acquire the metadata about the media files and obtain the in-

formation such as authors (file owner), file name, creation time and a link to the actual media into the database. The reader may note that techniques such as clustering [36] and Bayesian networks [35] can also be applied to this problem.

4.5 Event Storage

The implementation of the event model in *eVITAE* can be described through the entity-relationship diagram shown in Fig. 4. The key element in this diagram is the entity *Event* which, as its name suggests, corresponds to the key notion of the event-based conceptual model. It should be noted that the notion of an event here compresses the three-layered event model into a single layer. The *Event name* is a surrogate for a unique identifier that is generated for each event when it is created. An event is further described by a set of entities that are shown on the left side of Fig. 4. In the following, we briefly describe each of these entities in terms of the role they play in the event model:

Spatial Characterization: The entity *Space* is used to describe the spatial characteristics associated with events and is stored as latitude and longitude. The location information can be used in visualization and querying using map-based interfaces (see following section for details). Since, directly working with latitude-longitude information is cumbersome to most users, a look-up mechanism is created to map the latitude-longitude data to names of established places, such as cities along with their associated information like zip code, state and country. This helps users to interact with location-based information naturally and obviates the complexity associated with direct usage of latitude/longitude information

Representation of media: The entity *Media* is used to denote the media data which supports a given event. The media data may be referred to by a URL, a foreign key, or an index into a media-specific database (e.g. an image database). Further, additional information such as the media type may also be part of the description of this entity. It may be noted, that the model allows multiple media to support a given event as well as allowing a specific media to support different events.

From Information-Centric to Experiential Environments

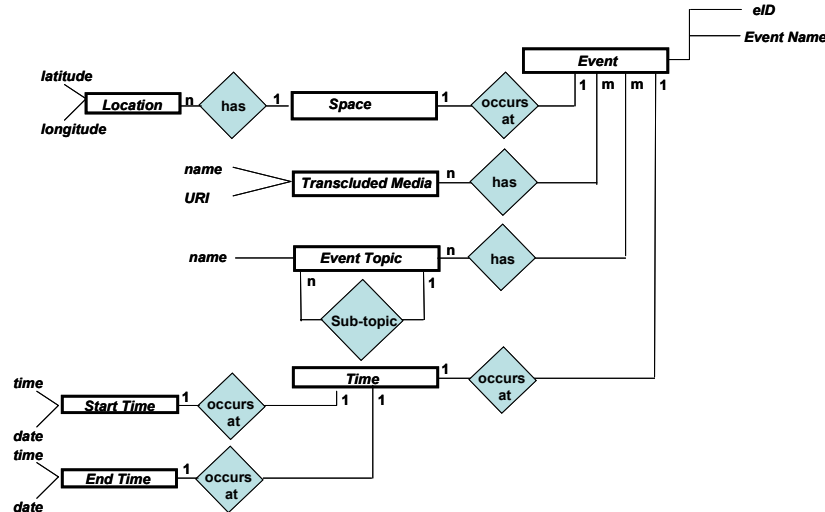


Fig. 4. ER-Diagram describing the event-model used for personal multimedia modeling

Representation of participants: In the context of personal media, social information, such as presence of an individual, can play an important role in information organization. This is emphasized by studies such as [32], where it was found that users associate their personal photographs with information on events, location, subject (defined as a combination of *who*, what, when, and where), and time. The entity *Participants* is used to depict such information. It allows, for instance, retrieval of all events (and associated media) where a specific person was present.

Temporal Characterization: The entity *Time* is used to model the temporal context of an event. Each event, in the ER-diagram is associated with a start-time and an end-time. In the case of point events, the start-time equals the end-time. It should be noted, that the physical implementation as described by the ER-model implicitly stores the *valid time* associated with events. This is because of the fact that in the contemporary setting many devices used for capturing personal information such as digital cameras (both for still photography and video capture), time-aware audio recorders, and electronic communications such as e-mails and instant messages allow direct and immediate information capture as an event occurs. However, if the domain semantics require keeping track of the *transaction time*, for instance, to have data available only for a specific time period after it has been published, the model can be extended in a manner similar to that used for valid time.

4.6 Event Presentation and Interaction Environment

In *eVITAE*, an integrated multimodal interaction environment is used as a unified presentation-browsing-querying interface. Two views of this environment are presented in Fig. 5 and depict its main components. The system employs direct manipulation techniques [26] to facilitate a natural user-system interaction. In this paradigm, users can directly perform different kinds of operations on items (events) of interest. Furthermore, combining the query and presentation space alleviates the cognitive load from the user's part unlike traditional query environments. Time and space are the primary attributes of the event definition, and hence are depicted as the primary exploration dimensions. Auxiliary panels are used to show the details of the events, and their attributes. Options for zooming, filtering, extraction, viewing relations, and details-on-demand are provided to help users interact with the information space. In the following, we discuss in greater details, the key aspects of the presentation and exploration environment:

Event representation: In the *eventCanvas* pane, events are represented through a recursive graphical representation called an *event-plane*. It consists of a rectangle which spans the duration of the event. The media, supporting an event, are represented by icons on the event-plane. Within an event-plane, the icons representing the media are chronologically ordered in terms of their capture times. The recursive nature of the event-plane is used to capture aggregate relationships where an event may comprise of other events. Such relationships, when they exist, are depicted using nested event planes. The primary purpose of such a representation is to provide users with a high-level view of the information that is independent of media specificities. When an event is selected, details about information associated with it such as the supporting media or alphanumeric attributes are automatically brought up in the *mediaDetailCanvas* and the *attributeCanvas* respectively. When a user needs to explore a specific event in detail, the media supporting an event is displayed in the *mediaDetailCanvas* by clicking on the event-plane icon of the corresponding event. When an event is thus selected, any available non-media attributes or alphanumeric meta-data related to it are simultaneously displayed in the *attributeCanvas*. Users have the option to see the actual media by clicking on a media icon in the *eventCanvas*. Selection of a specific media instance via clicking spawns a window which triggers the appropriate application for that particular kind of media.

From Information-Centric to Experiential Environments



Fig. 5. Two views of the event presentation and exploration environment in *eVITAE*: The screenshot on the left shows the overview of the information. The one on the right shows details of the event on spatial and temporal zoom-in. In both the views the top left pane is the *eventCanvas* where a chronologically ordered view of events is presented. The top right pane is used for visualization and interactions with the spatial aspects of the data. The bottom left pane is called the *mediaDetailCanvas*. Here, details of the specific media instances supporting an event are shown. Finally, the lower right pane, called the *attributeCanvas*, shows the non-media data, such as participant names or event descriptions, associated with a selected event.

Interactions with the temporal aspects of information: Events are defined over space and time. The *eVITAE* system not only captures this notion through event modeling, but also provides intuitive ways to visualize the spatio-temporal dimensions of the data. With respect to time, in the *eventCanvas* (see Fig. 5), a temporal distribution of events is presented with the events being ordered chronologically from left to right. A key operation supported in the *eventCanvas* is temporal zooming. Through it, users can zoom into a particular time interval to find more details about that time period (Fig. 5, right screenshot) or zooming out to see the overall distribution of events on the temporal axis. Support for local zooming (zooming within a specific interval) is also provided to allow focusing on a specific period for details, without the display getting cluttered by details over the entire timeline. Further, the semantic fisheye-view technique [29] is used to highlight the objects of current focus in the timeline while the user moves the slider across the timeline and zooms into a particular time interval. User can also select multiple intervals in the timeline, thereby creating multiple foci of the fish eye view.

Interactions with spatial information: Spatial information is displayed in the top-right panel of the *eVITAE* interface (see Fig. 5) and shows the overall distribution of events over space. The spatial display supports option to

Rahul Singh and Ramesh Jain

zoom down to a particular location by dragging a rectangle which contains that location (Fig. 5, right screenshot), and options to zoom out by clicking on the zoom out icon to get an overall picture of the information space. Furthermore, panning of the entire space is also supported. The spatial canvas in *eVITAE* has been implemented using an open source JavaBeans package called OpenMap [5]. A Mercator projection [39], in which meridians and parallels of latitude appear as lines crossing at right angles and areas farther from the equator appear larger, is used to display the various maps.

Interactions with alpha-numeric information: The presentation of alphanumeric information such as names of participants in an event or event descriptions is done using the *attributeCanvas*. Queries on events with respect to alphanumeric information can also be issued here. For example, to find all events having a specific participant a user would select the attribute “participants” in the *attributeCanvas* and type the name of the desired participant. The database is then queried for this information and the query results are displayed by highlighting the pertinent events in the *eventCanvas*.

Dynamic and Reflective User Interface: In a system having multiple simultaneous views of the data, such as *eVITAE* it is important to be able to establish relationships between different views of the dataset, such that any activity in one view is reflected in all the others. Such a capacity is essential for maintaining context as users interact with the information in different manners within each view. In *eVITAE* all the views of the data are tightly coupled through the database. For example, selecting an event in the timeline view leads to that event getting highlighted in the spatial view. Simultaneously, the details of that event are displayed in the *mediaDetailCanvas* and different attributes of the events are brought-up in the *attributeCanvas*. This in conjunction with support for rapid, incremental, and multi-modal interactions enables users to explore and “experience” the information from different perspectives.

4.7 Application 2: Business Activity Monitoring Application

Applications such as business activity monitoring (BAM) and homeland security must draw from a large network of disparate data sources, including databases, sensors, and systems in which data is entered manually. The goal of BAM is to allow a unified interface so that a manager can use it to monitor the status of activities at different locations and to analyze the causes of past events. In all such applications, real time data analysis must be combined with real time data assimilation from all sources to present a

From Information-Centric to Experiential Environments

unified model of the situation in an intuitive form. Techniques and tools developed for traditional database applications, such as payroll databases, are not adequate for this problem because a typical user is interested in exploratory formulations, such as understanding what could be the problem situations and why did they occur. In this context we note that data mining techniques are suitable when a hypothesis has been formed, but tools must first help in generating that hypothesis.

A cornerstone of our approach to this problem has been to create an environment that provides a holistic picture of all available information. By looking at the holistic picture, hypothesis can be formulated and then studied. Towards this we have developed an approach for implementing BAM systems that uses event-based domain model and the metadata to construct a new index that is independent of the data types and data sources. Specific event models have been developed for sales, inventory, and marketing domains. These models draw information from different databases, often from across the world, and unify this data around the domain events for each specific case. The reader may note that data in this problem does not display significant physical heterogeneity. However, the information assimilation challenges remain acute owing to the complexity and logical heterogeneity of the information.

All the events are stored in a database that is called *eventbase*. Similar to the problem of personal information management, the links to all original data or data sources are very important. These links are used to present appropriate media in the context of corresponding events. A strong interactive environment has been developed for users to interact with this system and gain insights through observations and analysis. The advantages of the approach include: (a) pre-processing important information related to events and objects based on domain knowledge, (b) presenting information using domain based visualization, and (c) providing a unified access to all information related to an event in terms of valid time. As an interactive environment for the system, an interface called the *EventViewer* has been developed which offers multidimensional navigational and exploration capabilities. An application screen of the *EventViewer* for the BAM application is shown in Fig. 6.



Fig. 6. Screenshot of an EventViewer for Demand Activity Monitoring

For an event three basic characteristics are its name and class, the location, and the time. As shown in the top left part of the screen in Fig. 6, a user can navigate through the class ontology hierarchy. Navigation through the location and time dimensions is either through zooming or by moving along different directions using simple natural controls. These traversals are very similar to those in video games. One can select parts of a map ranging from part of a room to the map of the world. Similarly, on the time line, one could be in the range of microseconds to centuries, or even larger scales when required. Once a user selects specific event classes, a region on the map, and a time interval, the system responds by presenting all events (and their selected attributes), satisfying these constraints. This information is presented by using the following three representations: (1) as a list in the space provided for event lists, (2) as symbols displayed on the location map, and (3) as symbols displayed at appropriate time points on the time line. These three representations are tightly linked. For instance, if an item in the list is selected, it gets simultaneously highlighted in the location and time displays. Such an approach to information search is a quintessential example of the WYSIWYG search philosophy.

From Information-Centric to Experiential Environments

A major goal of the BAM environment, as mentioned earlier, is to provide an intuitive feel of what events may have occurred and how they are related to other events and information. By presenting events on a map as well as on a time line, the context of the events is maintained and displayed to a user. A user can then refine the search criteria and as the criteria are refined, the results change appropriately. This instantaneous feedback allows users to experiment with the data set and develop insights and form hypotheses. When a user is interested in knowing more about a specific event, he or she can explore that event by double clicking on its representation in any of the three display areas. The system then provides all the data sources (like audio, video, or text) related to the selected event.

5. Conclusions

Traditional interface to database systems have been designed under the assumption that their role is to provide precise information as a response to precise queries. This model implicitly assumes that exact queries can be issued to obtain all relevant information. Therefore, such interfaces are not required to be interactive. Homogeneity in data, relatively small data volumes, and strongly structured application domains ensure the success of such information centric approaches. However, the volume and nature of the data being stored in databases today is significantly different than what was common a decade or more ago. Moreover, databases are being used for different roles now; the evolution in technology has put databases in the heart of systems where people seek not just data but information and insights. Dealing efficaciously with this new scenario requires query environments to become more exploratory and interactive.

Towards this, in this chapter, we have presented the idea of experiential environments for facilitating user-data interactions. In this paradigm, information is presented in a manner that seeks to take advantage of the sentient nature of human beings along with their cognitive and sensory abilities to experience and assimilate information. As examples, two experiential systems for supporting user interactions in different application contexts are presented. With the ever increasing availability of heterogeneous, media-rich data and requirements for supporting information assimilation across them, we believe, that the ideas of experiential environments will find active use in both research and development.

References

1. J. Allen. Maintaining Knowledge About Temporal Intervals. *Communications of the ACM*, Vol. 26, No. 11, 1983.
2. R. Belew. *Finding Out About: A Cognitive Perspective on Search Technology and the WWW*. Cambridge University Press, Cambridge, U.K., 2000.
3. D. Barreau. Context as a Factor in Personal Information Management Systems. *J. Am. Soc. For Information Science*, Vol. 46, No. 5:327-339, 1995.
4. "Log your Life Via Your Phone", March 10, 2004, <http://news.bbc.co.uk/2/hi/technology/>
5. BBN Technologies. *OpenMap - Open Systems Mapping Technology*, 1999
6. D. Barreau and B. Nardi. Finding and Reminding: File Organization from the Desktop SIGCHI Bulletin 27, No. 3:39-43, 1995.
7. T. Blum, D. Keislar, J. Wheaton, and E. Wold. Audio Databases with Content-based Retrieval. *Proc. IJCAI Workshop on Intelligent Multimedia Information Retrieval*, 1995.
8. M. Carey et al. Towards Heterogeneous Multimedia Information Systems: The Garlic Approach. *Fifth Intr. Workshop on Research Issues in Data Engineering -Distributed Object Management*, 1995.
9. C. Carson, M. Thomas, S. Belongie, J. Hallerstein, and J. Malik. Blobworld: A System for Region-Based Image Indexing and Retrieval, *Proc. Intr. Conf. on Visual Information Systems*, 1999.
10. K. Chakrabarti, K. Porkaew, and S. Mehrotra. Efficient Query Refinement in Multimedia Databases. *ICDE*, 2000.
11. L. Chen, M. Tamer Özsu, and V. Oria. Modeling Video Data for Content Based Queries: Extending the DISIMA Image Data Model. *MMM-2003*: 169-189, 2003.
12. N. Dimitrova et al. Applications of Video-Content Analysis and Retrieval. *IEEE MultiMedia* 9(3): 42-55, 2002.
13. J. D. N. Dionisio and A. Cardenas. A Unified Data Model for Representing Multimedia, Timeline, and Simulation Data", *IEEE Trans. On Knowledge and Data Engineering*, Vol. 10, No. 5, 1998.
14. S. Dumais, E. Cutrell, J. Cadiz, G. Jancke, R. Sarin, and D. Robbins. Stuff I've Seen: A System for Personal Information Retrieval and Re-Use, *ACM Conference on Research and Development in Information Retrieval*, 2003.
15. R. Elmasri and S. Navathe. *Fundamentals of Database Systems*, Addison-Wesley, 2004.
16. M. Flickner, H. Sawhney, W. Niblack, J. Ashley, Q. Huang, B. Dom, M. Gorkani, J. Hafner, D. Lee, D. Petkovic, D. Steele, and P. Yanker. Query by Image and Video Content: The QBIC System, *IEEE Computer*, 1995.
17. E. Freeman and S. Fertig. Lifestreams: Organizing your electronic life. In *AAAI Fall Symposium: AI Applications in Knowledge Navigation and Retrieval*, Cambridge, MA, 1995.
18. J. Gemmel, G. Bell, R. Lueder, S. Drucker, and C. Wong. MyLifeBits: fulfilling the Memex vision. *ACM Multimedia*, pp. 235-238, 2002.

From Information-Centric to Experiential Environments

19. S. Gibbs, C. Breiteneder, and D. Tschritzis. Data Modeling of Time-Based Media, Proc. SIGMOD, pp. 91-102, 1994.
20. W. I. Grosky and P. Stanchev. An Image Data Model. VISUAL 2000: pp. 14-25, 2000.
21. A. Gupta, T. Weymouth, and R. Jain. Semantic Queries with Pictures: The VIMSYS Model, Proc. 17th Intr. Conf. on Very Large Databases, 1991.
22. R. H. Gutting and M. Schneider. Realms: A Foundation for Spatial Data Types in Database Systems, Proc. 3rd Intr. Symp. On Large Spatial Databases, pp. 14-35, 1993.
23. R. H. Gutting. An Introduction to Spatial Database Systems, VLDB Journal, Vol.3, No. 4, 1994.
24. L. M. Haas, R. J. Miller, B. Niswonger, M. T. Roth, P. M. Schwarz, E. L. Wimmers. Transforming Heterogeneous Data with Database Middleware: Beyond Integration, IEEE Data Engineering Bulletin, 1997.
25. A. Y. Halevy. Answering Queries Using Views: A Survey, VLDB Journal, Vol. 10, No. 4, pp. 270-294, 2001.
26. E. L. Hutchins, J. D. Hollan, D. A. Norman. Direct Manipulation Interfaces, User Centered System Design. Lawrence Erlbaum Associates, 1986.
27. R. Hull. Managing Semantic Heterogeneity in Databases: A Theoretical Perspective, ACM Conf. on Principles of Database Systems, pp 51- 61, 1997.
28. R. Jain. Experiential Computing, Communications of the ACM, Vol. 46, No. 7, 2003.
29. P. Janecek and P. Pu. A Framework for Designing Fisheye Views to Support Multiple Semantic Contexts. International Conference on Advanced Visual Interfaces (AVI '02), ACM Press, 2002.
30. P. C. Kanellakis and D. Q. Goldin. Constraint Programming and Database Query Languages, Proc. 2nd Conference on Theoretical Aspects of Computer Science, 1994.
31. M. Lenzerini. Data Integration: A Theoretical Perspective, Proc. SIGMOD, pp. 233-246, 2002.
32. H. Libermann and H. Liu. Adaptive Linking Between Text and Photos Using Common Sense Reasoning, In Adaptive Hypermedia and Adaptive Web-Based Systems, deBra P, Brusilovsky P, and Conejo R (eds.), Springer-verlag, Berlin, pp. 2-11, 2002.
33. J. Li and M. Ozsu. STARS: A Spatial Attributes Retrieval System for Images and Video, Intr. Conf. on Multimedia Modeling, pp. 69-84, 1997.
34. T. Little and A. Ghafoor. Interval-Based Conceptual Models for Time-Dependent Multimedia Data. IEEE Trans. On Knowledge and Data Engineering, Vol. 5, No. 4, pp. 551-563, 1993.
35. J. Luo, E. Savakis, S. Etz, and A. Singhal. On the Application of Bayes Network to Semantic Understanding of Consumer Photographs, ICIP, 2000
36. A. Loui and A. Savakis. Automated Event Clustering and Quality Screening of Consumer Pictures for Digital Albuming. IEEE Trans. On Multimedia, Vol. 5, No. 3, 2000

Rahul Singh and Ramesh Jain

37. B. S. Manjunath and W. Y. Ma. Texture Features for Browsing and Retrieval of Image Data. *IEEE Trans. On Pattern Analysis and Machine Intelligence*, Vol 18, No 8, pp 837-842, 1996
38. H. Garcia-Molina, J. D. Ullman, J. Widom. *Database Systems. The Complete Book*, Prentice Hall, 2002
39. http://science.nasa.gov/Realtime/rocket_sci/orbmech/mercator.html
40. T. Nelson. Xanalogical Structure, Needed Now More Than Ever: Parallel Documents, Deep Links to Context, Deep Versioning, and Deep Re-Use”, *ACM Computing Surveys*, Vol. 31, 1999
41. J. Rekimoto. Time-machine computing: a time-centric approach for the information environment”, *Proceedings of the 12th annual ACM symposium on User interface software and technology*, pp. 45 – 54, 1999.
42. http://www.dcvIEWS.com/press/Ricoh_RDC-i700G.htm
43. S. Santini and A. Gupta. Principles of Schema Design for Multimedia Databases. *IEEE Trans. On Multimedia*, Vol. 4, No. 2, 2002.
44. S. Santini, A. Gupta, and R. Jain. Emergent Semantics Through Interaction In Image Databases, *IEEE Trans. On Knowledge and Data Engineering*, Vol. 13, No. 3, pp. 337-351, 2001.
45. S. Shekhar and S. Chawla. *Spatial Databases: A Tour*, Prentice Hall, 2003.
46. H-Y. Shum, M. Liao, S-F. Chang. *Advances in Multimedia Information Processing - PCM 2001*, Second IEEE Pacific Rim Conference on Multimedia, 2001.
47. R. Singh, Z. Li, P. Kim, D. Pack, and R. Jain. Event-Based Modeling and Processing of Digital Media, *Proc. First ACM SIGMOD Workshop on Computer Vision Meets Databases (CVDB)*, 2004.
48. R. Singh, R. L. Knickmeyer, P. Gupta, and R. Jain *Designing Experiential Environments for Management of Personal Multimedia*, *ACM Multimedia*, 2004.
49. H. Sridharan, H. Sundaram, and T. Rikakis. *Computational Models for Experiences in the Arts and Multimedia*, ETP, 2003.
50. D. Stan, I. K. Sethi. eID: a system for exploration of image databases. *Inf. Process. Manage.* 39(3): 335-361, 2003.
51. V. S. Subrahmanian. *Principles of Multimedia Database Systems*, Morgan Kaufman, 1998.
52. K. Toyama, R. Logan, A. Roseway, and P. Anandan. *Geographic Location Tags on Digital Images*, *ACM Multimedia*, pp. 156-166, 2003.
53. D. C. Tsichritzis and F. H. Lochovsky. *Data Models*, Prentics-Hall, New Jersey, 1982.
54. G. Widerhold, S. Jajodia, and W. Litwin. *Dealing with Granularity of Time in Temporal Databases CAiSE91*, pp. 124-140, 1991.
55. E. Wold, T. Blum, D. Keislar, and J. Wheaton. *Content-based Classification, Search and Retrieval of Audio IEEE Multimedia*, Vol 3, No. 3, pp 27-36, 1996.
56. M. J. Wynblatt. *Control Layer Primitives for the Layered Multimedia Data Model*, *ACM Multimedia*, 1995.