QBISM: A Prototype 3-D Medical Image Database System

*Manish Arya, William Cody, Christos Faloutsos**
IBM Almaden Research Center
San Jose, California

Joel Richardson
The Jackson Laboratory
Bar Harbor, Maine

Arthur Toga
UCLA
Medical School

1 Introduction

The goal of the QBISM[1] project is to study extensions of database technology that enable efficient, interactive exploration of numerous large spatial data sets from within a visualization environment. In our work we are currently focussing on the logical and physical database design issues to handle 3-dimensional spatial data sets created from 2-dimensional medical images.

Our specific application is the Functional Brain Mapping project at the Laboratory of Neuro Imaging, UCLA School of Medicine [12]. The goal of the brain mapping research is to discover spatial correlations between activity in the brain and functional behavior, e.g. speaking or arm movement. Such activity in the brain is frequently characterized by localized, non-uniform intensity distributions involving sections or layers of brain structures, rather than uniform distributions across complete structures. Discovering the precise locations of brain activity, correlating it with anatomy, and constructing functional brain atlases is the goal of an ongoing major medical research initiative. Ultimately, this understanding has clinical applications in diagnosis and treatment planning, as well as scientific and educational value.

To support the requirements of exploratory queries on multiple 3-D images, we have built an experimental prototype using the extensibility features of the Starburst DBMS developed at IBM's Almaden Research Center. The prototype has a client/server architecture, with IBM's Data Explorer/6000 visualization package serving as the foundation for an interactive, query front-end.

Section 2 describes the particular medical research problem we are studying and its query and data characteristics. Section 3 presents the status of our prototype. Section 4 summarizes the overall project and describes its future directions.

2 The Medical Application

2.1 Problem Definition

As mentioned, the purpose of QBISM is to support the data manipulation and visualization needs of the brain mapping project. The system we envision must support queries across multiple medical image studies in a very investigative, interactive, and iterative fashion. A study is actually a "billing" term referring to a set of medical images collected for a single purpose on a single patient, such as a 50 slice MRI[2] study or three x-rays of a fractured elbow. Querying and visualizing collections of

*On sabbatical from the University of Maryland at College Park. His work was partially supported by SRC and by the National Science Foundation (IRI-8958546), with matching funds from Empress Software Inc. and Thinking Machines Inc.

[1]Query By Interactive, Spatial Multimedia

[2]Magnetic Resonance Imaging. MRI images show soft-tissue structural information.

studies [3] will enable the return of statistical and comparative responses. Such capabilities will extend the power of medical visualization environments which today typically deal with a single study at a time. The following scenario, which is representative of the queries that medical researchers (i.e., those at the UCLA Laboratory of Neuro Imaging) would like to ask, illustrates a sample session with such a system in which each step generates a database query:

- The medical researcher may start by specifying and rendering a set of brain structures, for example those supporting the visual system, from a standard atlas [10].
- After repositioning the scene to a desired viewing angle, structures may be colored with a patient's PET[3] study data to highlight activity along their surfaces (see Figure 1(c)).
- The range of intensities in these structures may be histogram segmented, and then other regions in this PET study may be identified that have matching distributions.
- An arbitrary region in the study may be specified for visual comparison of its intensity pattern with the same or nearby region from a previous PET study.
- Paths for targeting electrodes or radiation beams that focus on an arbitrary region of interest may be calculated or simulated to permit the visualization of anatomical structures spatially intersected.
- An individual PET (or other study) may be statistically compared with data from a comparable subpopulation of the same demographic group to assess abnormality.

2.2 Data Characteristics

The database consists of a large collection of 3D studies and a set of anatomic atlases of the human brain.

Each study is a 3-dimensional scalar field (a 3-dimensional array of scalar values) representing some measured quantity, such as glucose consumption as an indicator of physiological activity, at each point in space. Studies are collected via an assortment of medical imaging modalities used to capture structural (e.g. MRI, CT[4], histology[5]) and functional / physiological (e.g. PET, SPECT[6]) information about the human brain. These studies typically consume about 1 - 30 megabytes, using current spatial resolutions and image depths, and could potentially consume over a hundred megabytes with increasing resolutions and depths. As a reference point, at The University of Virginia, which has a large medical center, the number of such tomographic studies can range from 5,000 to 15,000 per year, depending on modality. The total number of all types of imaging studies at the same medical center is 181,000 per year, resulting in 3 terabytes of data per year. Table 1 contains a breakdown for the tomographic studies. Modern hospitals are beginning to store all this imagery in systems known as PACS (Picture Archival and Communication Systems [13]). Using this data for advanced medical applications further increases the storage requirements due to the need to save derived data. Such data is a result of transformations to align and register the raw data, to create models suitable for volume and surface rendering of the data, and to build database representations that enable exploratory query.

As mentioned above, the database also contains atlases of reference brains, one for each demographic group. Each atlas models the exact shapes and positions of anatomical structures in the corresponding

[3] Positron Emission Tomography. PET images show physiological activity.
[4] Computed Tomography. CT images show hard-tissue structural information (e.g., bones).
[5] Histology images are acquired by physically slicing and photographing tissue, one thin layer at a time.
[6] Single Photon Emission Tomography. SPECT images, like PET images, show physiological activity.

Modality	Studies/Year	Images/Study	Image Size (bits)
CT	14810	30	512x512x12
MRI	5418	50	256x256x12
PET	6134	26	256x256x8

Table 1: Yearly tomographic study statistics for the University of Virginia Medical Center.

reference brain. A study itself does not identify the structure to which each voxel (3-D pixel) belongs, but an atlas can provide this information when overlaid on top of the study. Such use of an atlas is illustrated in the previous scenario by the second step, in which the brain structures of the visual system are used to retrieve parts of a particular patient's PET data.

Note that an acquired radiological study of a patient is not perfectly aligned with the corresponding atlas. Warping techniques [11] are used to derive affine transformations that allow a study to be registered to one or more appropriate atlases. In QBISM, we store the original study, the warped ones, and the warping transformations. The details of the warping are outside the scope of this paper. However, these automatic or semi-automatic warping algorithms are extremely important for this application. It is precisely this technology that permits anatomic structure-based access to acquired medical images as well as comparisons among studies, even of different patients, as long as they have been warped to the same atlas. Furthermore, it enables the database to grow, and be queryable, without time-consuming manual segmentation of the data.

3 A Prototype Implementation

We built a prototype that runs on IBM RISC System/6000 workstations. It integrates and utilizes the extensibility features of the Starburst relational DBMS [8] and the IBM Data Explorer/6000 (DX) scientific visualization product.

The user specifies a query by choosing a modality, a study, anatomic structures of interest, and intensity ranges of interest. The system then renders the selected data in 3D in one of several ways (see Figure 1). The user can manipulate the result in DX by changing the viewpoint, adding a cutting plane, or generating an isosurface, for example, or refine the query itself to select data from a different patient/study or part of space.

The system stores all large objects (e.g., studies and atlas structures) in Starburst long fields [4] and the associated attributes in relations. New SQL functions we added to Starburst perform the necessary spatial operations, such as "intersection()". A new processing module we added to DX accepts the user's query and communicates with Starburst through a network connection to retrieve the spatially-restricted answer. Details on the data types and their representations, the operations, and performance experiments are in [2].

4 Summary and Future Directions

We have discussed the requirements and the initial implementation of QBISM, a prototype system for managing and visualizing 3D medical images. In order to allow convenient querying over multiple studies, we believe that such a system should be built on top of an extensible DBMS engine, appropriately extended to handle spatial data, and combined with a high-quality visualization tool as the user interface. The challenges in the project are to define and implement the database extensions that

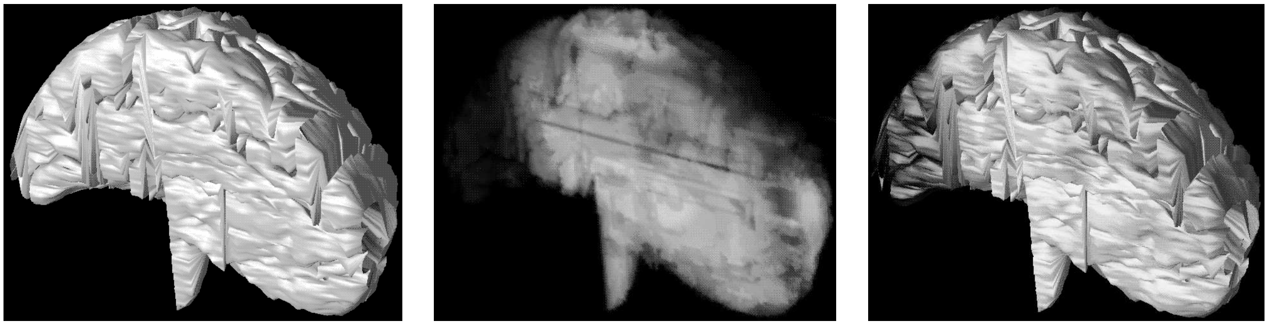

(a) (b) (c)

Figure 1: Sample query results. (a) One brain hemisphere from the atlas. (b) The intensity data from a PET study inside the hemisphere. (c) The same PET data mapped onto the surface of the hemisphere. Note the difference in shading between a and c, which is more prominent in color.

support medical researchers' ad-hoc queries over populations of studies with interactive response times, despite the large size of even a single study.

Future goals of QBISM include:

- Development of multi-study indexing techniques for anatomic atlases and patient studies to accelerate queries over large patient populations.
- Addition of approximate spatial representations and the associated filter/refinement query processing strategies to further optimize large-scale queries [7].
- Incorporation of data mining and hypothesis testing techniques to support investigative queries. An example of a hypothesis testing query is *"is it true that people with dyslexia show this intensity range in this region of their PET studies?"* An example of a data-mining/rule-discovery query is *"find PET study intensity patterns that are associated with any known neurological condition in any subpopulation"*. Data mining algorithms for relational databases are presented in [1].
- Integration of support for query by image content. We would like to support similarity queries like *"find all the PET studies of 40-year old females with intensities inside the cerebellum similar to Ms. Smith's latest PET study"*, or sub-pattern matching queries, like *"find patients whose MRI studies show a hippocampus with shape similar to that of this patient who has a particular neurological disorder"*. Clearly, we need feature extraction and similarity measures. Research in machine vision has yielded several good features for 2-D images, e.g. the "QBIC" project at IBM ARC [6] proposed and experimented with some color, shape and texture features. Our challenge is to discover appropriate features for 3-D medical images.
- Development of natural, spatial interaction mechanisms to help pose queries with 3-D regions of interest and to help manipulate the results. Effective solutions may require special hardware, such as a 3-D mouse or a "data glove".

Acknowledgments: We would like to thank Walid Aref and Brian Scassellati for helping initiate this work; the Starburst developers for providing advice and help with Starburst; and the UCLA LONI Lab staff for providing and helping interpret the human brain data.

References

[1] R. Agrawal, T. Imielinski, A. Swami, "Mining Association Rules between Sets of Items in Large Databases", ACM SIGMOD, 1993 (to appear).

[2] M. Arya, W. Cody, C. Faloutsos, J. Richardson and A. Toga, "QBISM: Extending a DBMS to Support 3D Medical Images", submitted for publication, 1993.

[3] H. Fuchs, M. Levoy and S. M. Pizer, "Interactive Visualization of 3D Medical Data", IEEE Computer, 22, 8, pp. 46-51, Aug. 1989.

[4] T.J. Lehman, B.G. Lindsay, "The Starburst Long Field Manager", Procedings of the 15th International Conference on Very Large Data Bases, Amsterdam, pp. 375-383, Aug. 1989.

[5] A.D. Narasimhalu and S. Christodoulakis "Multimedia Information Systems: The Unfolding of a Reality", IEEE Computer, 24, 10, pp. 6-8, Oct. 1991.

[6] W. Niblack, R. Barber, W. Equitz, M. Flickner, E. Glasman, D. Petkovic, and P. Yanker, "The QBIC Project: Querying Images By Content Using Color, Texture, and Shape", SPIE 1993 International Symposium on Electronic Imaging: Science & Technology, Conference 1908, Storage and Retrieval for Image and Video Databases, February 1993.

[7] J. Orenstein "Redundancy in Spatial Databases" Proc. of ACM SIGMOD conf., Portland, Oregon, May 1989.

[8] P. Schwarz, W. Chang, J.C. Freytag, G. Lohman, J. McPherson, C. Mohan, and H. Pirahesh, "Extensibility in the Starburst Database System," Proc. 1986 Int'l Workshop on Object-Oriented Database Systems, Pacific Grove, September 1986, pp. 85-92.

[9] A. Silberschatz, M. Stonebraker, J. Ullman (eds.), "Database Systems: Achievements and Opportunities", Communications of the ACM, Vol. 34, No. 10, pp. 110-120, Oct. 1991.

[10] J. Talairach and P. Tournoux, "Co-planar stereotactic atlas of the human brain", Thieme, Stuttgart, 1988.

[11] A.W. Toga, P. Banerjee and B.A. Payne, "Brain warping and averaging", Int. Symp. on Cereb. Blood Flow and Metab., Miami, FL 1991.

[12] A.W. Toga, "A digital three-dimensional atlas of structure/function relationships", J. Chem. Neuroanat., 4(5):313-318, 1991.

[13] A.W.K. Wong, R.K. Taira and H.K. Huang "Digital Archive Center: Implementation for a Radiology Department", AJR 159, pp. 1101-1105, November 1992.