



In the NEWS

Features Editor: Dennis Taylor
dtaylor@computer.org

ALSO FEATURED

LIFE ANNOTATION: STORING AND SEARCHING OUR PERSONAL DIGITIZED MEMORIES

Toward More Intelligent Healthcare

Danna Voth

With AI's help, researchers are creating tools to help doctors make better diagnoses, gather enhanced information from medical tests, improve their surgery and examination skills, and more quickly make relevant biomedical discoveries.

Improving diagnoses

Getting a quicker, smarter look at medical images taken using computer-aided tomography (CAT scans) and magnetic resonance imaging would help doctors make faster, more accurate diagnoses. At Simon Fraser University, Ghassan Hamarneh and Chris McIntosh have developed virtual worms that can “crawl” through images of tubular structures such as blood vessels, airways, and spinal cords. These 3D crawlers analyze the image data using AI reasoning techniques and present their interpretation on a screen.

Initially, Hamarneh analyzed medical images by inserting a deformable geometrical model into the image. This crawler deformed when it was attracted to certain structures in the image. But, says Harmeneh, the deformable models weren't very reliable.

So, he and his team extended the model by giving the crawler perception intelligence (the ability to interpret or identify structures that are sensed) and decision-making capabilities. A “vesselness filter” helps the improved crawler know when it's in a tubular

structure. The filter collects image intensity data and uses it to determine whether the brightness level indicates a tubular structure—something to crawl down, such as a blood vessel.

The crawler's feature-detecting algorithms are based on a Hessian matrix that formally describes how the intensity changes. The algorithms obtain both eigenvalues and the corresponding eigenvectors. “The nice thing is the smallest eigenvector of this matrix points along tubular structures,” Hamarneh says.

To detect branches in the tubes, the crawler has a spherical sensor that also collects image intensity data. The presence of multiple areas with a distinct brightness intensity, highlighted using image-processing algorithms, could indicate branching. Decision-making algorithms help the crawler decide where to crawl and when to stop. The crawler decides to follow a branch when it detects a bifurcation and stops when the tube narrows to the point that the tubular structure is ending.

Hamarneh and his team are working on

making the GUI more user friendly and making the crawler faster by porting their prototype system, which uses the MATLAB programming language, to ITK (the Insight Segmentation and Registration Toolkit), which uses C++. Next, Hamarneh says, they will tackle the more complex problem of interaction between multiple crawlers.

Improving surgery planning

Vipin Chaudhary and his team at the University at Buffalo are harnessing the brute force of high-performance computers to analyze and predict organ and tumor positions during operations. They've been working on methods for predicting the trajectory of the smallest possible incision that will remove only the tumor and do the least damage to the surrounding tissue. In neurosurgery, Chaudhary says, “when you make an incision, the cerebrospinal fluid flows out, and due to gravity most of the time the brain sinks.” The problem is that this shift can be as big as, or even larger than, the tumor. “Because most of these surgeries are image guided, the images you have were taken before the incision was made,” Chaudhary says. “The surgeon now has to look at the image and predict in his mind where the different structures would be.”

Chaudhary's system does that for the surgeons. Employing pattern recognition, it correlates images of the patient's brain to a medical atlas. It then optimizes the trajectory for the surgeon's incision. As the doctor operates, the system takes data from the doctor's probe about where the structures are in the brain and provides the new

IEEE Computer Society Publications Office

10662 Los Vaqueros Circle, PO Box 3014
Los Alamitos, CA 90720-1314

STAFF

Lead Editor

Dennis Taylor

dtaylor@computer.org

Group Managing Editor

Crystal R. Shif

cshif@computer.org

Senior Editors

Dale Strok, Shani Murray, and Linda World

Assistant Editor

Brooke Miner

Editorial Assistant

Molly Mraz

Publication Coordinator

Hilda Carman

Contributing Editor

Joan Taylor

Production Editor

Jennie Zhu

Technical Illustrations

Alex Torres

Publisher

Angela Burgess

aburgess@computer.org

Associate Publisher

Dick Price

Membership/Circulation Marketing Manager

Georgann Carter

Business Development Manager

Sandra Brown

Senior Production Coordinator

Marian Anderson

Submissions: For detailed instructions and formatting, see the author guidelines at www.computer.org/intelligent/author.htm or log onto *IEEE Intelligent Systems'* author center at Manuscript Central (www.computer.org/mc/intelligent/author.htm). Visit www.computer.org/intelligent for editorial guidelines.

Editorial: Unless otherwise stated, bylined articles as well as products and services reflect the author's or firm's opinion; inclusion does not necessarily constitute endorsement by the IEEE Computer Society or the IEEE.

position information, helping the doctor adjust his or her movements as the structures change position.

Chaudhary has started a company, which has built a prototype that will be tested this spring at the Detroit Medical Center. Chaudhary plans to extend the research to orthopedic and pediatric surgery.

Helpful feedback for surgeons

At the Johns Hopkins Whiting School of Engineering, Gregory Hager and his colleagues are formulating a surgery grammar that helps break down repeated motions into observable components, much like speech comprises phonemes. Early research suggested to Hager and his team that they could think of surgical motions as gestures consisting of smaller subunits they call "gestemes." They trained hidden Markov models to recognize specific gestemes for tasks such as membrane peeling in retinal surgery.

Using models of surgeries captured on video by the da Vinci surgical robot, Hager's team next looked at the skill of suturing. "We would sit down and analyze the video and say at this point they were pulling the suture, and at this point they were handing the suture from left hand to right hand, and so on and so forth," Hager says. Then they applied statistical learning methods, such as linear discriminant analysis, hidden Markov models, support vector machines, and Bayes classifiers, to their data to develop a statistical classifier that can replicate their classification of the data. They achieved 92 to 93 percent reliability. More recently, they have been building Gaussian mixture models, which have increased the reliability to 95 to 97 percent. Hager hopes to build a prototype using the gestemes to help surgeons improve their skills by offering specific critiques of their hand motions.

Reading colonoscopies

A research team is developing a database management system for colonoscopy videos that can help objectively measure those exams' quality. The team members are Johnny Wong and Wallapak Tavanapong at Iowa State University, JungHwan Oh at the University of North Texas, and Piet C. de Groen at the Mayo Clinic College of Medicine.

Their system has three major components. The automatic-capturing component captures the video stream, forwards the data to the analysis server, and isolates the video frames of single procedures.

The automatic-quality-analysis component analyzes videos of a procedure and generates objective measurements for that procedure. This component employs algorithms that determine

- how far the colonoscopy probes have traveled, on the basis of accumulated forward movements of the instrument,
- whether an image is clear or blurred, and
- whether all the mucosa (the moist tissue lining the digestive tract) was examined.

The algorithms use machine learning methods such as clustering techniques and classifiers. This analysis is currently a post-procedure process, but the team hopes to develop a real-time version to help endoscopists during a live exam.

Finally, the reporting-system component produces a report based on the analysis results.

The team has started a company to develop, sell, and support computer-aided quality control systems for endoscopy. Their prototype is being tested at the Mayo Clinic Rochester, and they plan other clinical trials to evaluate the system. They hope to have their software in use at clinics by 2008.

Gathering nuggets of medical knowledge

The probability of following and understanding all the data from recent medical studies is understandably low—millions of medical articles are published, too many for one person to read and process. So, Chitta Baral and Graciela Gonzalez at Arizona State University have developed a way to make the intellects of many people available to everyone. Their computer program CBioC (Collaborative Bio Curation) can, with a little help from individuals, analyze and organize the enormous collections of biomedical data available in journals.

Baral says that recently, "post docs in biology would be hired to read all the papers, find nuggets of knowledge, and put it in a database." That's very expensive, and he says, "the medical companies who do that keep it to themselves because they invest so much money in doing that." Another method is to employ natural language processing to extract important facts from the papers. However, Baral says the NLP programs make a lot of errors.

So, Baral and Gonzales developed a system that relies on reader annotations. The

system first performs automatic extraction using natural language processing. It then asks readers to vote on whether they think the information is relevant and accurate. They can exclude wrong extractions or add new ones.

To make sure that the collaborative part works, the team added a trust management component. The component's algorithms assign the value 1 to all initial votes. As people vote for or against an extraction, the weighting shifts. For example, if out of 10 votes, eight are yes and two are no, the yes votes will receive a greater weight. The weighting also considers how many times a particular voter has voted against the grain—perhaps to confuse the system.

However, Baral and Gonzales haven't yet

gotten the number of users they need to make the algorithms work. (You can download the software for free at www.cbioic.org.) "We have hundreds of votes," Baral says, "but for this to really work, it has to get to the thousands."

Intelligent automation of clinical examinations and surgical-skills coaching could help doctors learn skills more quickly and improve their effectiveness, especially when this includes the sharing of best practices. The search to find the best way to practice medicine will also benefit from the synergy resulting from collaboration between intelligent systems and humans.

Life Annotation: Storing and Searching Our Personal Digitized Memories

Sara Reese Hedberg

Personal computers are no longer simply work tools. They now store the intermingled memories of our personal and professional lives in gigabytes of photos,

emails, calendars, documents, videos, and so on. A small cadre of researchers is beginning to tap these digital stores through life annotation, enabling us to search our own PCs the way we search the Web.

As we may remember

The vision for life annotation stems from Vannevar Bush, a top US scientific and military policy maker during World War II and the Cold War. In the 1945 *Atlantic Monthly* article "As We May Think" (www.theatlantic.com/doc/194507/bush), Bush outlined a device called a "memex" that would be "an enlarged intimate supplement" to human memory.

The problem that Bush was addressing remains the same: "The summation of human experience is being expanded at a prodigious rate, and the means we use for

threading through the consequent maze to the momentarily important item [has not kept pace]."

Today we have vast, cheap digital storage—roughly US\$1 per gigabyte. The question is, how do we find what we're looking for? Drawing from memory research in psychology, neuroscience, AI, and computer science, researchers at Microsoft, the University of Southampton, the Massachusetts Institute of Technology, and elsewhere are stepping up to the memex blueprint Bush posited more than 60 years ago.

Gordon Bell's MyLifeBits

In 1999, the legendary Gordon Bell, who led the development of the Digital Equipment Corporation's revolutionary VAX minicomputers, was inspired by Bush's article (and a challenge from AI researcher Raj Reddy) to digitize his life's work and memories. Now at Microsoft Research, Bell has since captured 160 Gbytes of his papers, books, presentations, photos, videos, files, and so on in MyLifeBits (<http://research.microsoft.com/barc/>

**Intelligent
IEEE Systems**

How to Reach Us

Writers

For detailed information on submitting articles, write for our Editorial Guidelines (isystems@computer.org) or access www.computer.org/intelligent/author.htm.

Letters to the Editor

Send letters to

Dennis Taylor, Lead Editor
IEEE Intelligent Systems
10662 Los Vaqueros Circle
Los Alamitos, CA 90720
dtaylor@computer.org

Please provide an email address or daytime phone number with your letter.

On the Web

Access www.computer.org/intelligent for information about *IEEE Intelligent Systems*.

Subscription Change of Address

Send change-of-address requests for magazine subscriptions to address.change@ieee.org. Be sure to specify *IEEE Intelligent Systems*.

Membership Change of Address

Send change-of-address requests for the membership directory to directory.updates@computer.org.

Missing or Damaged Copies

If you are missing an issue or you received a damaged copy, contact membership@computer.org.

Reprints of Articles

For price information or to order reprints, email isystems@computer.org or fax +1 714 821 4010.

Reprint Permission

To obtain permission to reprint an article, contact William Hagen, IEEE Copyrights and Trademarks Manager, at copyrights@ieee.org.

MediaPresence/MyLifeBits.aspx). He has even stored personal memorabilia such as photos of his mother's birth certificate, coffee cups, plaques, and genetic information. He adds 1 Gbyte each month, including photos from the miniature research camera, SenseCam, that he sports around his neck. "[MyLifeBits] reduces the clutter of physical information," Bell explains. It's also useful. "I recently had to introduce an important person," he says. "I had scanned my calendars and knew the date and place of a 1983 meeting that [resulted in] a significantly more personal introduction."

In 2000, fellow researcher Jim Gemmell began building a software infrastructure to unify the fragmented data of Bell's growing MyLifeBits into a collected corpus. His team used a potpourri of C#, SQL, and other code. The shell is now to the point where Microsoft Research has funded a handful of related university projects that use MyLifeBits software and/or the SenseCam. At Columbia University, for instance, researchers are applying AI techniques for content analysis to automatically segment and index audio files (<http://labrosa.ee.columbia.edu/projects/personalaudio>); Dublin City University has a similar project for image files (www.cdvpc.dcu.ie/SenseCam).

The MyLifeBits software is still experi-

mental, and the search and retrieval mechanisms can come up short. To address this, Bell and Gemmell are considering incorporating some of the AI-based capabilities that fellow Microsoft researchers in Eric Horvitz's group have built.

Microsoft's LifeBrowser

Recently Horvitz (<http://research.microsoft.com/~horvitz>) demonstrated his Adaptive Systems and Interaction Group's LifeBrowser system. As he turned on his monitor to begin, Horvitz beamed, "This is the holodeck of my life!"

At almost warp speed, he showed where he was on election day, including pictures of whom he was with and what he did during and after work. He showed a list of documents he "touched" on another day, noting the color coding for those he had actually edited versus read. He also requested a short slide show of his "best" family pictures from the 4th of July several years ago. An artistic, dynamic montage of varied images quickly appeared—not cached, but chosen and rendered on the fly.

Although on the surface, LifeBrowser makes searching and serving content look effortless, powerful intelligent software lies underneath, with a well-thought-out interface on top. The design is based on studies of

human memory indicating that humans often use special events or "landmarks" to guide recall. These can be public events such as 9/11 or personal events such as family trips.

The system automatically predicts landmark events. A calendar crawler working with Microsoft Outlook extracts multiple properties from calendar events (for example, location, organizer, and relationships between participants). The system then uses Bayesian machine learning and reasoning to automatically derive atypical features from events—because it's often the unusual, not the quotidian, that's memorable.

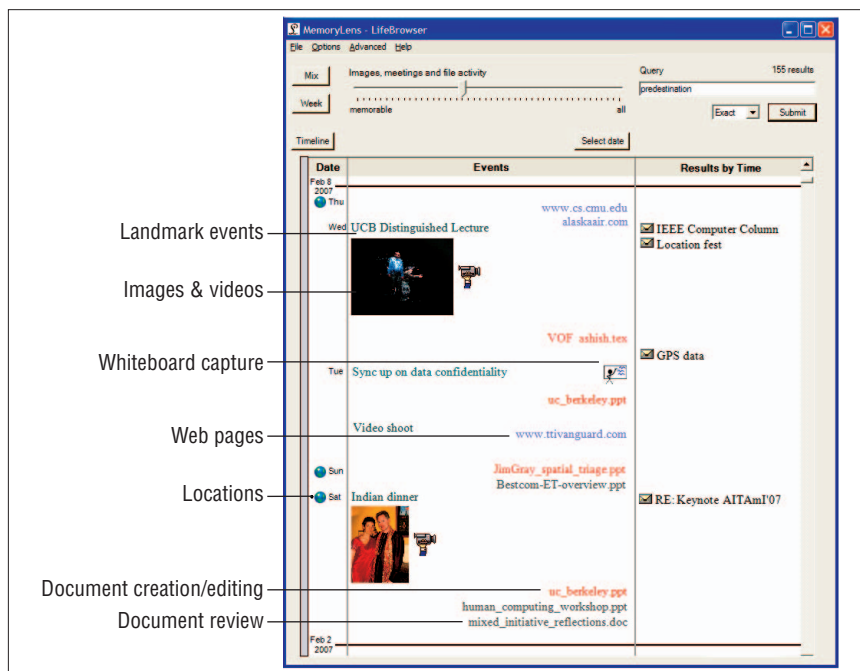
Because images also help humans remember, an image crawler can analyze a photo library. Using information stored by a digital camera, and photo features automatically extracted by image analysis algorithms, the system employs Bayesian learning to predict the most likely landmark pictures, such as those selected for Horvitz's election day diary and his 4th of July slideshow.

"It's interesting how much a part of my life [LifeBrowser] has become," observes Horvitz. "At work, I find myself often using the trace of my activity structured along a timeline to quickly find documents, presentations, and Web pages. It's very natural to find exactly what I am looking for by skimming among a mix of key memory milestones. Broadening beyond work, my family now views most of the images and videos we capture—without doing any of our own sorting or annotation."

Horvitz and some of his group members regularly use LifeBrowser. It's also scheduled for large-scale distribution to thousands of Microsoft employees.

Researchers describe life annotation as "life changing," "a security blanket for memory," and "freeing the brain for more creative pursuits." The applications are broad—from a digital diary for posterity and progeny, to a prosthetic for patients with memory loss. Some even envision the day when all our health data, behaviors, and habits are recorded and provided to insurance companies.

Clearly, life annotation involves legal, ethical, and political issues. "Society is going to have to deal with [these issues]," notes Bell. "We aren't necessarily introducing any more problems. We expose them by making a system that is probably easier to exploit." ■



LifeBrowser view of an automatically constructed timeline of landmark events and activities.