

Developing a public software resource as a shared foundation for future medical image research.

OPEN SOURCE SOFTWARE FOR MEDICAL IMAGE PROCESSING AND VISUALIZATION

By **TERRY S. YOO** and **MICHAEL J. ACKERMAN**

Societies often create smaller subsets or communities that connect with one another for commerce and intellectual exchange over mutual interests. In science and engineering, the need for communication among researchers is often hampered by artificial barriers of university politics, economic market forces, and the sheer momentum of an academic reward structure that values individual discovery over joint development. Recent initiatives have attempted to reduce some of these barriers, encouraging collaborative multidisciplinary research programs. Through this effort, we have studied the processes that lead to the successful foundation of new communities.

Our current focus has led to the creation of Insight, a project for open source image processing software development, along with the Insight Software Consortium (ISC), which includes more than 17 participating universities and commercial institutions. The initial emphasis of this effort is to provide public software tools for 3D segmentation and deformable and rigid registration, capable of analyzing the head and neck anatomy of the Visible Human Project data. The eventual goal is to provide the cornerstone of a self-sustaining software community in 3D, 4D, and higher-dimensional data analysis. Ultimately, we intend this to be a public software resource that will serve as a foundation for future medical image research.

The Visible Human Project was initially formed to collect data from human subjects to serve as a guidebook and baseline data set in modern anatomy research and education [1]. The intent of the current effort is to amplify the investment being made through the Visible Human Project and future programs for medical image analysis by reducing the reinvention of basic algorithms. We are also hoping to empower young researchers and small research laboratories with the kernel of an image-analysis system in the public domain. We are committed to open source public software, including open interfaces supporting

connections to a broad range of visualization and graphical user-interface platforms.

Community through Scientific Rendezvous


There have been several notable examples of accelerating computer science research through what we describe as a scientific rendezvous. Generally, a scientific rendezvous is a meeting of scientific minds: shared data, a shared conference, or more essentially a shared vision that promotes dialogue throughout a community. It is most likely to be an open and welcoming group, admitting free dialogue across many subjects, but since it primarily advances a small emerging but significant field, the topics will be narrow in scope and tightly focused.

A scientific rendezvous is not likely to be closed, proprietary, nor a standard. Some of the most important scientific rendezvous in computer science have been open source software projects that have stimulated education, research, and training. The creation of the Unix operating system along with the long-term support of the University of California at Berkeley did more to disseminate advanced operating system design elements among academic and research institutions than any particular text or curriculum [5]. The open hardware architecture of the IBM PC

OPEN SOURCE SUPPORT SOFTWARE

The Insight Project has led to the development of some powerful and important independent software tools. The need for concurrent development of source code for Windows, Macintosh, and Unix-based systems has led to the creation of CMake, a cross-platform software build tool that creates both makefiles for Unix systems as well as project files to support users of Microsoft and Borland compilers. In addition to CMake, the ISC has created DART, a build environment that uses a dashboard, an online resource for viewing the instantaneous state of the build process. ITK is built automatically and continuously as changes are checked into the source code repository. By convention, all Insight software must be accompanied by build and regression tests. This requirement permits the immediate testing of the software after the build is complete, and the results of the tests, including compile failures, simple runtime errors, and warnings are posted on a Web-based table. The ITK Dashboard is one of the essential tools we use to coordinate the many geographically distributed development teams.

The ISC has adopted generic programming, a form of typeless programming through the aggressive use of C++ templates in our software development. This type of programming allows us to manipulate images without worrying about the atomic type of the image (pixel or voxel) elements. However, the wide use of templates makes wrapping of the code for multiple language binding difficult. The solution was provided through Cable, a method that instantiates all necessary templates, parses the resultant C++ source code, and creates wrapper code permitting the calling of ITK routines from other languages. Currently, Cable supports TCL (Tool Control Language), Python, and Java.

These software engineering ideas, CMake, DART, and Cable, have been some of the important by-products of this initiative. The support tools are also all open source, and are being employed in similar projects across an international community of software teams [12]. Together, they may constitute as valuable a contribution to software systems development as the other project goals are to medical imaging. 

also introduced a common basis over which a community dialogue could be conducted. The Apple Macintosh graphical user interface did much to codify conventions in user-interface design, leading to a new jargon, a protracted discussion on pull-down vs. pop-up menus, and the promotion of a wide-ranging commercial effort in mouse and trackball design.

Scientific rendezvous sometimes emerge from funded research projects rather than commercial ventures. Perhaps the most successful sponsored project that led to the formation of a community dialogue and later to an entire industry was the development of Internet tools. HTML serves as the core for the most successful of these tools, and the open nature of HTML opened the way for derivative products such as Netscape Communicator and Internet Explorer Web browsers. The National Library of Medicine (NLM) has participated in multiple efforts of this nature, including GenBank, the public gene sequence repository. The Visible Human Project has produced a pair of studies of human anatomy that have provided common ground from which to extend multidisciplinary research in the fields of anatomy, medical imaging, visualization, and other areas of computer science [1].

What distinguishes a successful scientific rendezvous? The fundamental strength is its ability to promote and sustain dialogue for whatever field it supports. It provides a foundation for greater research across institutions, demanding interoperability rather than proprietary functions. As such it should encourage communication and stimulate development, often through a new lexicon of technical terms that form around the emerging technology. This concept is based on widely adopted conventions rather than formal standards.

The demand for common ground usually leads to the development of standards; we submit that before a field is mature enough to require standards, it can benefit from some nucleus around which scientific and engineering conversations can revolve. Principal attributes of the most successful scientific rendezvous in IT are open architecture (software or hardware), extensible infrastructure and control languages, and explicit support for the substance of the core material including distribution mechanisms, open licensing, tutorials and documentation, tools, and bug fixes.

The Need for Public Image Analysis Tools

There are pressing needs for automated techniques in image analysis. Medical processes such as early cancer detection, monitoring the progress of medical treatment, response to cardiovascular disease, and

analysis of neurological disorders such as stroke all benefit from advanced imaging. Public exchange of software tools accelerates the development and improvement of medical care through imaging. This need creates a focus for the community and a basis for a scientific rendezvous.

After conducting a joint National Institutes of Health (NIH) workshop in 2000 [2], the NLM, in a multipartner arrangement, supported the creation of the Insight Toolkit (ITK), an application programming interface for high-dimension image processing. This work has been the focus over which both research and government institutions have come together.¹

Planning the Consortium

Most of the systems, conferences, or data sets that have become the common basis for information exchange have evolved spontaneously as need and opportunity converged. They seldom represent good business models. Indeed, the idea of making the components open and accessible limits individual control of any element, inhibits the exclusive generation of profit, and invites competition. Anyone planning a potential scientific rendezvous should attempt to include as many of the elements of previous successful projects and recognize that neither control nor profit are useful motives or goals in such an endeavor.

We sought a consortium of software developers, rather than a single award, in the hopes that the natural process of building consensus would generate a broad versatile software base for a wider audience. Since "design by committee" efforts are often unsuccessful, we empowered a core group of software architects and systems integrators to make primary design choices. The ISC is working to deliver a software toolkit to improve and enable research in volume imaging for all areas of health care. The successful collaboration among such disparate groups as engineering schools, medical schools, small businesses, and large corporate R&D centers participating in the project is a measure of the success of this effort.²

Design Principles. The project evolved along certain principles, eventually reinforced and extended during the first meeting of the ISC in November

¹Members of the funding coalition include the National Institute for Dental and Craniofacial Research, the National Eye Institute, the National Institute of Mental Health, the National Science Foundation, the National Institute for Neurological Disorders and Stroke, the National Institute on Deafness and other Communication Disorders, the DoD Telemedicine and Advanced Technology Research Center, and the National Cancer Institute.

²This group continues to grow internationally, demonstrating increasing momentum for this concept and proving its value. ITK public software developers include participants from: NLM, Carnegie Mellon, Cognitica, Columbia, GE Global Research, Georgetown, Harvard Brigham and Women's Hospital, Imperial College London, Insightful, Kitware, Mayo Clinic, UPenn, University of Pittsburgh, and the University of Utah.

1999. The fundamental guiding principals include:

- Open source (see the sidebar “Open Source Software Support”);
- No cost licensing;
- Portable (runs on Windows and Unix systems);
- Compact (encumbers no cascades of licenses);
- Designed to be easily maintained;
- Emphasis on 3D; and
- Validation metrics included as part of the tools.

Target Uses. Several target uses for the toolkit were identified during the planning stages, including:

- Repository of image-processing algorithms;
- Standard implementations for comparison;
- Prototype methods to encourage variations;
- A common format for writing new methods;
- Communication for imaging software research;

ALGORITHMS AVAILABLE IN ITK

ITK has representative implementations of a wide variety of classification, segmentation, and registration filters as shown in the figure here. These methods include segmentation, classification, and registration filters [3, 8], as well as medial modeling primitives for analysis and measurement [10]. Mesh operations are also included in ITK along with an entire finite-element modeling system [4]. Linear and nonlinear image filters (including nonlinear scalar and vector-valued diffusion) as well as level-set methods [6] are all part of ITK. Some Insight teams have worked on hybrid segmentation tools, combining fuzzy-connectedness concepts with deformable finite-element meshes, initialized with a Voronoi image decomposition [9]. A high-level overview of the wide cross-section of available methods includes:

- Software foundation for young researchers; and
- Metrics with which to compare methods.

Algorithms and Validation

The Insight Project was formed around the need for public medical image segmentation and registration software, including the archiving of working versions of often-cited common algorithms enabling comparison of existing techniques with improved variations. ITK also emphasizes threaded, pipelined, and streaming implementations suitable for execution in multiprocessor environments and for use on large data sets that do not fit within core memory.

Template-based generic programming has permitted the flexible combination of solvers, optimizers, and other methods with typical approaches for image segmentation and registration. For example, a statistical affine normalized correlation image registration tool can be implemented with gradient descent opti-

Image Processing

- Linear and edge-preserving smoothing
- Edge detection and feature extraction
- Arithmetic and morphological algorithms
- Support for vector-valued data


Segmentation

- Statistical classification/clustering
- Connected components and fuzzy connectedness region growing
- Hierarchical watersheds
- Level-set methods and deformable models
- Hybrid algorithms

Registration

- Generalized framework with interchangeable metrics, transforms, and optimizers
- Cross-modality registration
- Multiresolution support
- PDE and FEM-based deformable registration

Examples of medical image segmentation and registration algorithms available in ITK: a) Functional MRI fused with MR angiography using landmark initialized mutual information registration (courtesy of UNC); b) 2D to 3D registration of angiography data (courtesy of the Imperial College of London); c) Inner-ear segmentation of the cochlea and vestibular system using fast-marching level-set methods (courtesy of Kitware); d) Eyes, muscles, and optic nerves of the Visible Human Project data using interactive color watershed segmentation (courtesy of the University of Utah).



Some of the most important scientific rendezvous in computer science have been OPEN SOURCE SOFTWARE PROJECTS THAT HAVE STIMULATED EDUCATION, RESEARCH, AND TRAINING.

mizer and a regular-step gradient descent optimizer with a simple reorganization of existing templates, simply by recombining software building blocks. Thus, a single technique may be represented multiple times in the toolkit through significant variations. Meshes and images represent the core classes of ITK, with filters used as operations. A wide variety of algorithms are represented in ITK (see the sidebar “Algorithms Available in ITK”).

In order to measure the efficacy of a segmentation or registration method, data with known truth must be available [7]. Performance of an algorithm is also measured along many different dimensions including efficiency, accuracy, and precision [11]. Team members are working to publish test metrics as well as public data, including the software framework necessary to execute their planned validation pilot studies. The Insight Project is designed with the concept of academic transparency at its heart. All implementations are available in a working form, and all performance reports will be accompanied by the tests performed to measure them. Researchers following in the wake of the original implementers are encouraged to use the same tests to evaluate the performance of their improved or alternate segmentation or registration methods. Furthermore, all software developers in this area are strongly encouraged to contribute to this growing corpus of work in image processing.

Conclusion

A non-profit association, the ISC was incorporated in early 2004, and is an indication of the future of this effort. In time, the ISC may become the trustees of the open source software and serve as a focal point for the community. At the time this article was written in late 2004, the project had produced its fifth release of ITK (1.8). NLM is committed to supporting this effort, and regular software releases are intended. The software, which is publicly available in source-code form, builds and runs on Windows systems as well as on a variety of Unix systems including Solaris, Linux, SGI Irix, and Macintosh OS-X, all on a wide variety of compilers; interested software developers should visit the Insight Web site: www.itk.org. An open mailing list, insight-users, is currently active and open to all subscribers. Support tools including CMake, Cable, and DART, the dashboard support infrastructure, are all part of the public software associated with this project and available as source code.

Time, performance, and the broad user community will determine the ultimate success of the Insight project as a scientific rendezvous. To date, it has been a rare privilege to observe and participate in the growth and evolution of this noble test of group dynamics and collaborative software development. ■

REFERENCES

1. Ackerman, M.J. The Visible Human Project. In *Proceedings of the IEEE* 86, 3 (1998), 504–511.
2. Ackerman, M.J., Yoo, T.S., and Jenkins, D. 2000. The Visible Human Project: From data to knowledge. In H.U. Lemke et al., Eds., *Computer Assisted Radiology and Surgery (Proceedings of CARS2000)*. Elsevier, Amsterdam (2000), 11–16.
3. Chalana, V., Ng, L., Rystrom, L., Gee, J., Haynor, D. Validation of a brain segmentation and tissue classification algorithm for T1-weighted MR Images. In K.M. Hanson and M. Sonka, Eds., *Proceedings of the SPIE Medical Imaging Conference* (2001), 1873–1882.
4. Chen, T. and Metaxas, D. Gibbs Prior models, Marching Cubes, and Deformable Models: A hybrid framework for 3D medical image segmentation. In *Proceedings of the MICCAI 2003 Conference* (Montreal, CA, Nov. 2003), 703–710.
5. Collinson, P. Unix: The cult. In *Proceedings of Usenix* (1986), 22–28.
6. Elangovan, V. and Whitaker, R. From surfaces to sinograms: A direct approach to the segmentation of tomographic data. In *Proceedings of the MICCAI International Conference* (2001), 213–223.
7. Gee, J. Performance evaluation of medical image processing algorithms. In K.M. Hanson, Ed., *Proceedings of the SPIE Medical Imaging Conference* (2000), 19–27.
8. Ibanez, L. and Bullitt, E. Synthetic clusters as a 2D/3D registration technique for image-guided spinal surgery. In *Proceedings of the Computer Assisted Orthopedic Surgery Conference* (Pittsburgh, PA, 2001).
9. Imielinska, C., Metaxas, D., Udupa, J., Jin, Y. and Chen, T. Hybrid segmentation of anatomical data. In *Proceedings of the 4th MICCAI International Conference* (2001).
10. Stetten, G. and Pizer, S. Medial node models to identify and measure objects in real-time 3D echocardiography. *IEEE Transactions on Medical Imaging* 18, 10 (1999), 1025–1034.
11. Udupa, J.K. and Herman, G.T., Eds. *3D Imaging in Medicine*. CRC, Boca Raton, FL, 2000.
12. Yoo, T., Ackerman, M.J., Lorensen, W., Schroeder, W., Chalana, V., Aylward, S., Metaxas, D., Whitaker, R. Engineering and algorithm design for an image processing API: A technical report on ITK—The Insight Toolkit. In J. Westwood, et al. Eds., *Medicine Meets Virtual Reality*, IOS Press, Amsterdam (2002), 582–592.

TERRY S. YOO (tyoo@mail.nih.gov) is head of the 3D Informatics Program at the National Library of Medicine at the National Institutes of Health in Bethesda, MD.

MICHAEL J. ACKERMAN (ackerman@nlm.nih.gov) is the assistant director of High Performance Computing and Communications at the National Library of Medicine at the National Institutes of Health in Bethesda, MD.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.