



Published in final edited form as:

Med Image Anal. 2016 October ; 33: 176–180. doi:10.1016/j.media.2016.06.035.

Increasing the Impact of Medical Image Computing Using Community-Based Open-Access Hackathons: the NA-MIC and 3D Slicer Experience

Tina Kapur, Ph.D.¹, Steve Pieper, PhD², Andriy Fedorov, PhD¹, J-C Fillion-Robin, MS³, Michael Halle, PhD¹, Lauren O'Donnell, PhD¹, Andras Lasso, PhD⁴, Tamas Ungi, MD, PhD⁴, Csaba Pinter, MS⁴, Julien Finet, MS³, Sonia Pujol, PhD¹, Jayender Jagadeesan, PhD¹, Junichi Tokuda, PhD¹, Isaiah Norton, MS¹, Raul San Jose Estepar, PhD¹, David Gering, PhD⁵, Hugo J.W.L. Aerts, PhD¹, Marianna Jakab, M.S.¹, Nobuhiko Hata, PhD¹, Luiz Ibanez, PhD⁶, Daniel Blezek, PhD⁷, Jim Miller, PhD⁸, Stephen Aylward, PhD³, W. Eric L Grimson, PhD⁹, Gabor Fichtinger, PhD⁴, William M Wells, PhD¹, William E. Lorensen, MS⁸, Will Schroeder, PhD³, and Ron Kikinis, MD¹

¹Brigham and Women's Hospital and Harvard Medical School

²Isomics Inc.

³Kitware Inc.

⁴Queens University

⁵Healthmyne Inc.

⁶Google

⁷Mayo Clinic

⁸General Electric Research

⁹MIT

Abstract

The National Alliance for Medical Image Computing (NA-MIC) was launched in 2004 with the goal of investigating and developing an open source software infrastructure for the extraction of information and knowledge from medical images using computational methods. Several leading research and engineering groups participated in this effort that was funded by the US National Institutes of Health through a variety of infrastructure grants. This effort transformed 3D Slicer from an internal, Boston-based, academic research software application into a professionally maintained, robust, open source platform with an international leadership and developer and user communities. Critical improvements to the widely used underlying open source libraries and tools -- VTK, ITK, CMake, CDash, DCMTK -- were an additional consequence of this effort. This project has contributed to close to a thousand peer-reviewed publications and a growing portfolio of US and international funded efforts expanding the use of these tools in new medical computing

applications every year. In this editorial, we discuss what we believe are gaps in the way medical image computing is pursued today; how a well-executed research platform can enable discovery, innovation and reproducible science (“Open Science”); and how our quest to build such a software platform has evolved into a productive and rewarding social engineering exercise in building an open-access community with a shared vision.

Keywords

open access; medical image computing; reproducible research; 3D Slicer; hackathon; project week; NA-MIC; open source; open science

Notable Successes

Traditionally and understandably, reason for the gap between theory and practice is ascribed to the academic reward system and the mechanisms for research funding where novelty is valued over robustness and reproducibility. Nevertheless, several groups have overcome the funding odds and cultural barriers to create tools that have stimulated community building and open science. The specific choices made by different groups differ in terms of target application areas, license, architecture, sustainability and maintenance strategies, and each tool has tradeoffs that should be considered before incorporating it as the basis for development. Notable and widely used examples include *ImageJ*, *Analyze*, *Osirix*, *ClearCanvas*, *FreeSurfer*, *FSL*, *MITK*, *medINRIA*, *NifTK GIMIAS*, *SPM*, *MeVisLab*, and *3D Slicer*. *ImageJ* (NIH) is extensible and widely used for image viewing, but with little native support for 3D processing. *OsiriX* (UCLA, Pixmeo) is a popular Mac OS X radiology workstation with its source code freely available under GPL license, and an FDA-approved counterpart available to license commercially. *ClearCanvas* (University of Toronto, Synaptive Medical) began as both a free GPL-licensed DICOM viewer for Windows, and a commercially available product. More recently, work on the open source *ClearCanvas* has been discontinued. *Analyze* (Mayo Clinic) is perhaps the oldest commercial package for medical image analysis and visualization dating back nearly three decades and is also available commercially. *FreeSurfer* (MGH, UCSD) specializes in automatic parcellations of a human brain from MRI images. *FMRIB Software Library* (*FSL*, Oxford University) is a popular collection of image analysis and statistical tools for the analysis of functional, structural and diffusion MRI brain imaging data, is freely available for non-commercial use. *Medical Imaging Interaction Toolkit* (*MITK*, DKFZ Heidelberg) framework focuses on interactive applications (eg. in image-guided therapy) and has a well-considered infrastructure for the manipulation of objects in 2D and 3D views. It is freely available for non-commercial and commercial use. *medINRIA* (INRIA, France) notably provides capabilities for diffusion MR processing and tractography, and requires authorization by developers for further distribution of the software. *GIMIAS* (University of Sheffield, UK) is an environment for rapid prototypes development that focuses on computational physiology modeling, and is freely available for all uses. *SPM* (UCL, London) is a Matlab (Mathworks Inc.) toolbox that is heavily used to organize and interpret functional neuroimaging data. It is available under an LGPL license, which requires that all modified and extended versions of the program to be made free as well. *MeVisLab* is a commercial platform for developers,

enabling fast prototyping and translation into regulated commercial environments. Unlike the commercial or GPL-licensed software, VTK (the Visualization Toolkit) and ITK (Insight Segmentation and Registration Toolkit) are two very successful toolkits available freely for both commercial and research use. VTK was enhanced for a decade with diverse funding from the high performance computing (HPC) community, and it found widespread use in multiple research and commercial systems including those highlighted above; however, its base medical image computing capabilities were relatively unfunded and stagnant until it received a recent software maintenance grant from the NIH to support medical uses of VTK. ITK, a pioneer of community-based governance, has been recently re-energized with an architectural refresh after ten years. We are increasingly convinced that the widespread adoption of ITK has helped accelerate the capabilities and the level of complexity of MIC algorithms over the past ten years. In 2009, a transatlantic effort was formally launched for an open source “Common Toolkit” that is governed by a BSD-style license. The charter of CTK is to work on topics that are not covered by existing toolkits, and members of NA-MIC are active participants in it.

3D Slicer and the National Alliance for Medical Image Computing

3D Slicer is a free and open source software package (BSD-style license) for image analysis and scientific visualization building on our judgement of the best tools and practices available in the community, such as VTK (BSD license) and ITK (Apache 2.0 license). It is used in a variety of MIC research applications, including autism, multiple sclerosis, systemic lupus erythematosus, huntington’s disease, schizophrenia, neurosurgery, traumatic brain injury, orthopedic biomechanics, chronic obstructive pulmonary disease, lung cancer, breast cancer, cardiovascular disease, prostate cancer, and gynecologic cancer. 3D Slicer can be extended to enable development of both interactive and batch processing tools for a variety of applications.

Many, if not all, of the above examples of successful open source software, including 3D Slicer, started out as byproducts of research projects intended to investigate novel algorithms, publish papers, and perhaps create minimal prototypes as proofs of concept. Almost without exception, the founding leaders “repurposed” effort from other research funding to build infrastructure that was crucial for the productivity of their research team and allowed them to build efficiently upon previous work.

For 3D Slicer in particular, we started this process in 1997 when we envisioned the application of algorithmic advances from the Artificial Intelligence Laboratory (now part of the CSAIL) at MIT to translational research in neurosurgery and neuroscience at Brigham and Women’s Hospital. In 2004, we launched the National Alliance for Medical Image Computing (NA-MIC), which was officially mandated to build an open software platform for reproducible science. We had to “repurpose” no more! The US National Institutes of Health provided significant funding and several leading research and engineering groups joined us. 3D Slicer until that time was a Boston-based academic software application with capabilities for segmentation, registration, and visualization. It was written by MIT graduate students and used for research by Brigham and Women’s Hospital neurosurgeons. The NA-MIC effort transformed 3D Slicer into a professionally maintained, robust, and extensible

software platform for translational research by adopting high-quality software engineering practices and building upon and contributing back to the widely-used open-source libraries that form its foundation: VTK, ITK, CMake, and CDash.

Today, 3D Slicer continues to be free and open source, and it has many features that make it valuable for an active international community of users: “one-click” installers for multiple platforms (Windows, Mac and Linux); industry-strength engineering based on test-driven software processes; extensible with a community driven “App Store” (called the Extension Manager) that implements dozens of solutions for MIC problems; and support and training that is available online and at workshops throughout the year. An important point that we believe contributes to its popularity is that 3D Slicer is governed under a BSD-style license which essentially is a statement that anyone can do anything with the software, and the authors are not liable for any damage. The main advantage of this BSD-style license with respect to the GPL or LGPL licenses is that it makes Slicer-based code more palatable to companies for commercial product development: individual modules or the entirety of Slicer can be used as the basis for commercial products, integrated with proprietary technologies, and perhaps become the basis of a profitable business, without the imposition of royalties, restrictions, or special permissions, and without even involving 3D Slicer developers.

Organizing the Community for Growth

The community responsible for the success of 3D Slicer includes clinician scientists with knowledge of the clinical problems and image data that need analysis, computer scientists and physicists with novel algorithms based on elegant applied mathematics, software engineers with the ability to understand the clinical problems, data, and algorithms and then create reliable and useable tools, and biomedical engineers with the multidisciplinary skills to adapt and deploy these tools in a range of clinical settings. As the funding and membership of the community began to grow in 2004, so too did the list of deliverables that we promised the funding agencies.

We began organizing for this growth and launched a community event in 2005 that was inspired by the experiences of some of us at General Electric (GE). One inspiration was a six-sigma process improvement method called “work-out,” that was used successfully throughout the healthcare business division of GE (Ulrich et al. 2002). A workout is a structured forum to bring people together to solve problems as a team. By design, it engages producers more than managers, though the support of managers is critical. A workout starts by setting a specific, measurable goal, identifying the participants who need to be involved, and collecting relevant data prior to the workout event. Planning takes weeks, while the event itself is a few days long. The event begins with a discussion with the full group then participants work in small teams on specific subtasks, and come together again at the end to report results. A second inspiration for the hackathon was a unique environment in the early 2000s at the corporate research headquarters of General Electric, where 6 or more of us often did group programming to solve issues at a large screen with multiple back projectors. A strong sense of community and shared ownership of outcomes resulted from both the workouts and the group-programming sessions, and this spirit is what we sought for the growing 3D Slicer and NA-MIC community.

To the best of our knowledge, we were among the first in the field of medical image computing to implement such a forum, and now, with the experience of 22 “NA-MIC Project Week” events behind us, it makes perfect sense that variants of such events, popularly known as hackathons (or hackfests) are effectively used across disciplines as a means of channeling creativity towards a common mission (Briscoe & Mulligan 2014).

NA-MIC Project Week

The Pilot Hackathon

The first NA-MIC Project Week, held in the summer of 2005 at MIT, registered 44 attendees working in 15 teams. Every attendee was funded by the same NIH grant (or was a named collaborator on that grant), everyone funded by the grant was required to attend, and each project was connected to a deliverable of the grant. In retrospect, the arrangement has a closed and provincial feel to it, but since that grant included participants from 14 institutions, many of whom only knew a few of the other participants, it was an appropriate starting point. A notable outcome of this event was the adoption of a BSD-style license by 3D Slicer, a decision, which has had tremendous positive impact on its adoption in the years since.

The 22nd Hackathon

The 22nd Project Week, held in January 2016 at MIT, registered 77 participants working in 47 teams. Even though the original grant that provided funding to launch the Project Week series concluded in 2014 after a planned 10-year duration, these 47 teams represented 20 new funded efforts from around the globe. Each of the new teams was equally committed to the principles of open science as the original, and better informed about the impact of licenses on fluid exchange of even open and free software. Several projects were focused on building foundational components such that they can be leveraged by multiple funded efforts.

Best Practices

In the 11 years between the 1st and the 22nd Project Week events, we have had some time to hone what we believe are key ingredients for their successful organization. For the readers who may be considering organizing their community in a similar manner, we summarize these practices as well as insights that we think might be of value.

- *Frequency.* The event is held twice a year, which turns out to be just the right frequency for face-to-face meetings of this geographically widespread group. Several participants routinely attend only once a year, and others fit in additional smaller team meetings.
- *Physical co-Location.* The definition of participation is to physically show up at the venue. Of course, electronic communication including video-conferencing routinely supplements these events but is not yet considered a satisfactory substitute.

- *Duration.* These run a week: from Monday lunch to Friday lunch. We have experimented with shorter events and there is unanimous preference for this length.
- *Size.* Size ranges between 60–120 participants, and a good balance is achieved when ~20% are content experts or *gurus* whose primary role is to help, and 20% are first timers who infuse new ideas and projects into the mix.
- *Venue.* It is important for the event venue to provide a large room that can accommodate lecture-style seating for 100 people for 2 hours each at the start and end of the event, and be converted into a banquet-style arrangement with 25 round tables with 5 chairs each for the rest of the week; a second room that fits 20 people for breakout sessions; and good internet connectivity. Our events have all been held in major cities -- MIT campus in Boston, hotels in Salt Lake City, conference halls in Barcelona and Heidelberg -- where local organizers secured the logistics.
- *Fee and Food.* All attendees, including the organizers pay a registration fee. This fee is calculated to break-even with the cost of food and venue rental charges (when held at hotels or conference centers). Good food and an abundance of coffee are non-negotiable ingredients of such an event, and charging this registration fee allows us to scale this event easily.
- *Preparation.* Year round planning is carried out by the leaders to introduce interested people -- students, researchers, industry leaders -- to the concepts of open source medical image computing and this community in particular, with a recommendation to they join the mailing list (~700 members) when they become interested in attending one. Active planning for a specific event starts 6–8 weeks in advance with an email to the mailing list containing a schedule of weekly preparatory conference calls leading up to the event. In addition to these weekly conference calls which average ~10 attendees each, several individual conversations take place between leaders and prospective attendees. The goal of these conversations is for the leaders are to understand the goals and skills of each prospective participant, and to match them with a project or a team for the event. This step is perhaps the most important quality-control step in the organization of Project Week; participants with poorly defined goals can undermine the overall quality of the event.
- *Project Selection.* Projects that typically benefit the most are where: one participant seeks to essentially duplicate functionality that has been created at another institution; individuals from multiple institutions comes together to divide up a large task that is needed by all, but no one is sufficiently resourced to tackle it alone; one participant is gathering requirements to build a component that is needed by a group.

- *Open and Collaborative Agenda.* The agenda for the meeting is created collectively on a public wiki (wiki.na-mic.org), during the preparatory conference calls. For every event, the wiki page contains also a list of all projects (goals, teams, outcomes).
- *The Event.* The actual event begins on a Monday with lunch and a 2-hour introductory session where one person from each team uses their wiki page description to introduce the goals and members of the project in 90 seconds to the other participants. After this, laptops are plugged in and work is done in these teams with experienced members moving fluidly between 4 or 5 teams. 2 or 3 optional breakout sessions are held on topics of common interest (as determined during preparatory conference calls). We conclude on Friday with a 2-hour session in which each team reports progress against the goals that were set for the week, and records it on the wiki.
- *Leadership.* The leadership style for Project Weeks is much like that of a teacher in a flipped classroom (Skirpan & Yeh 2015); materials and guidelines are provided ahead of time, and face-to-face time is used for interactive, peer-driven learning that is steered by the leader.
- *Open is Key.* A key reason that the Project Week style of community organizing works for NA-MIC and 3D Slicer is because the underlying software is open and there are no barriers to sharing. We welcome those who take our software, enhance it, and either contributes the results back or make the result proprietary and commercial. Even if, for business reasons, some choose not to acknowledge that they have “3D Slicer inside” their products, we consider ourselves indirect beneficiaries of their success in part because taxes on their sales help fund government grants, but more importantly because new therapies will never be routinely applied to help patients unless they are embedded in FDA cleared commercial products.



Project Weeks have proven to be a highly successful model for community involvement, rapid progress, and for building a congenial community where new members can rapidly come up to speed on 3D Slicer usage and development. Since 2005, 22 of these open-access, open-source, extreme programming hackathons have taken place; 1928 participants (612 unique) from around the world have worked on 1098 projects to create a thriving open-access community for medical image computing.

About Leadership, Scope and Longevity—As the 3D Slicer software and community have grown, diverse trends and traits of successful research tool builders have become apparent. They are differentiated by leadership style, scope and longevity. We define **Level 1** projects as those that are linked to a single person. The strength of Level 1 projects is that they are typically very clean implementations of a cohesive vision and their weakness is that their scope and longevity is tied to one person. **Level 2** projects are multi-year, sustained efforts linked to a single group or institution. Implementations are less efficient than Level 1 because more people are involved, but in return there is a relative increase in scope and longevity. **Level 3** projects are decades-long community-based efforts that are led by a shared vision. Level 3 projects can tackle the largest scope, and have longevity beyond that of a single generation of leadership. The price to pay is the significant effort that is needed to maintain clear channels of communication within the community, and to ensure software processes support a larger community. 3D Slicer transitioned from a Boston-based Level 2 project to an internationally run Level 3 project in the last decade. The community around it is distributed around the globe: the users who download it over a thousand times per week, the 50+ developers who contributed code to it during the last year, the core leadership that spans Boston, Kingston (Ontario), Albany, and Chapel Hill, and the diverse portfolio of funded application projects. The funding is no small part of this Level 3 status; In 2004, there were only a handful of active grants that hinged on 3D Slicer; today there are more than 20. People who started out as 3D Slicer users in 2004 have graduated into funded investigators who are innovating with Slicer and contributing finished applications to the Slicer Extension Manager.

Into the Future

3D Slicer depends on a vibrant and active community, a strong leadership team, and dedicated outreach and support practices in order to survive as a Level 3 platform and application. This requires the software to adjust and respond to evolving concepts and capabilities of the underlying toolkits and libraries. Maintaining stability for the existing community of developers and users is critical to protect the existing investments, but not taking advantage of the new capabilities carries the risk of lagging behind the field. The strong leadership team that is integral to the 3D Slicer community must diligently review new capabilities and assess how to optimally integrate them into the architecture. Assessment of when an infrastructure is mature enough for incorporation into widely distributed software is one critical aspect of this ongoing process. We continue to pursue funding opportunities, like NA-MIC, that will advance the MIC fields by directly supporting infrastructure and community development.

In Closing

In summary of our position about the importance of open and reproducible science, we quote observations attributed to Jon Claerbout of Stanford (Buckheit & Donoho 1995) which are equally relevant today: “*An article about computational science in a scientific publication about the scholarship itself, it is merely **advertising** of the scholarship. The actual scholarship is the complete software development environment and the complete set of instructions which generated the figures.*” We are aligned with this vision of scholarship, and it is our hope that someday, every article published in MedIA will be accompanied by source code, data, and the “secret sauce” (Collins & Tabak 2014) of parameter settings that were used to generate the results. We believe that this Open Science approach will enable us all to do better science and ultimately provide better patient care.

Acknowledgments

This work was enabled by NIH Grants: U54EB005149, P41EB015902, P41EB015898, U24CA180918, U01CA199459, U24CA194354, R01EB020667, R01CA111288, U01CA19023, U24CA194354, R01EB014955, R01EB021396, R01HL11693, R01HL116473, and funding from Cancer Care Ontario.

References

- Briscoe, G.; Mulligan, C. Digital innovation: The hackathon phenomenon. London: Creativeworks London Work Paper; 2014. Available at: <http://www.creativeworkslondon.org.uk/wp-content/uploads/2013/11/Digital-Innovation-The-Hackathon-Phenomenon1.pdf>
- Buckheit, JB.; Donoho, DL. Lecture Notes in Statistics. 1995. WaveLab and Reproducible Research; p. 55-81.
- Collins FS, Tabak LA. Policy: NIH plans to enhance reproducibility. Nature. 2014; 505(7485):612–613. [PubMed: 24482835]
- Skirpan, M.; Yeh, T. Proceedings of the 46th ACM Technical Symposium on Computer Science Education. SIGCSE '15. New York, NY, USA: ACM; 2015. Beyond the Flipped Classroom: Learning by Doing Through Challenges and Hack-a-thons; p. 212-217.
- Ulrich, D.; Kerr, S.; Ashkenas, R. The GE Work-Out: How to Implement GE's Revolutionary Method for Busting Bureaucracy & Attacking Organizational Proble. McGraw Hill Professional; 2002.

Highlights

Traditional academic reward system and mechanisms for research funding create “valley of death” for research prototypes by valuing novelty over robustness and reproducibility.

The US-based National Alliance for Medical Image Computing (NA-MIC) was formed in 2004 to address this problem by creating robust tools that enable reproducible science.

NA-MIC transformed 3D Slicer, a research prototype for medical image processing, into a professionally maintained, robust, and extensible software platform for translational research that is downloaded over a thousand times a week by researchers worldwide.

A key ingredient in the success of NA-MIC is the strong sense of belonging to the community and a shared ownership of outcomes. This has been fostered through a semi-annual “NA-MIC Project Week” hackathon series that has run continuously since 2005.

In our view, the *medical image computing* or “MIC” community has been quite successful at developing innovative algorithms but less successful at building usable tools. There is a rich literature around extraction of information and knowledge from medical images using computational methods; a query in Google Scholar about “medical AND image AND (segmentation OR registration)” yields a count of over a million publications in the last decade. There is no doubt that as a field the MIC community has systematically invented algorithms and created compelling prototypes to aid medical discovery, diagnosis, and therapy monitoring in an environment with ever-increasing data and complexity. These results have been systematically published in highly regarded venues and ensured academic promotions for scientists in the field. What has been pursued less systematically is the creation of efficient pathways for these prototypes to become robust tools that can be used by other researchers and, when appropriate, fully commercialized. As a result, much of the MIC literature describes innovations that end up in the so-called “valley of death” where promising techniques never leave the prototype stage.

Here we discuss how our particular style of “Open Science” -- a transparent process to create fully reproducible and translatable methods, with open code, data, and tutorials licensed to support and encourage translation into clinical products -- has allowed us and others to create quality tools without sacrificing academic recognition and funding.

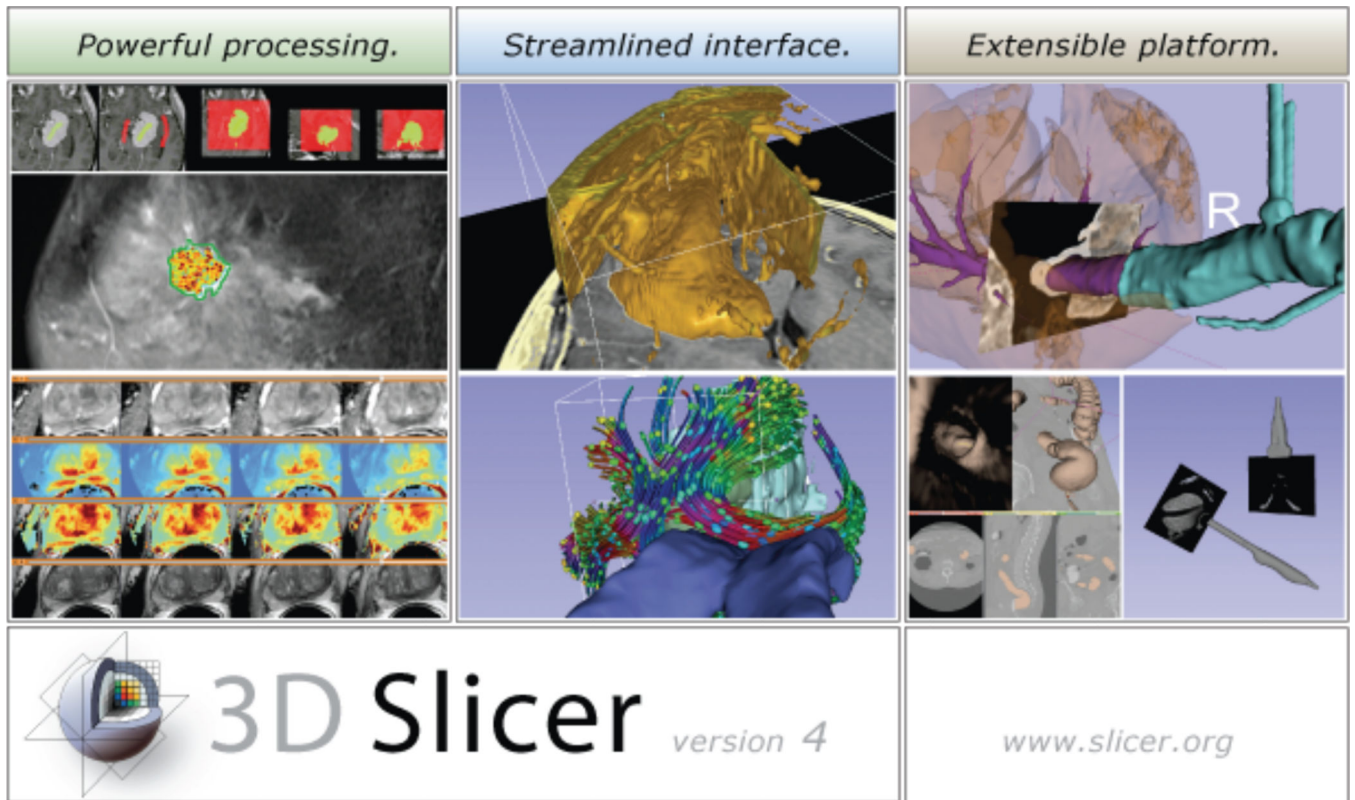


Figure 1. Representative applications of 3D Slicer (clockwise from upper left) in mammography, neurosurgery, image-guided therapy, colonoscopy, diffusion tractography, and prostate imaging.