SPECIFIC AIMS

Imaging has emerged as one of the key tools used by biomedical investigators to further our understanding of human biology in health and disease. With the growing importance of imaging modalities, research imaging programs have transitioned into high throughput enterprises serving investigators with a diverse range of skills and backgrounds. Imaging has also evolved into a highly quantitative enterprise requiring significant computational resources. As a result of the growing scope, scale, and sophistication of research imaging, investigators have come to rely on imaging informatics tools to assist in managing, processing, and sharing imaging and related data. Amongst these tools, **XNAT has emerged as the most widely used informatics platform across a wide range of imaging research**, including in cancer, neuroscience, and cardiology. XNAT is currently in use at dozens of institutions, supporting hundreds of independent studies and thousands of research personnel. Here we propose to continue the ongoing development of XNAT, its integration with other scientific software, and our efforts to support its use in the community.

The central function of XNAT is to manage a workflow for importing, archiving, processing, and securely distributing imaging and associated data. While this workflow has been at the core of XNAT since its inception, **the XNAT workflow is continuing to evolve to enable four key use cases**: *institutional image repositories*, which serve as central data management resources for multiple investigators and research studies; *clinical research hubs*, which provide a bridge between clinical imaging and research; *multi-center trial databases*, which require rigorous data oversight and quality assurance procedures; and *data sharing services*, which broadcast data selectively to collaborators or openly to the research community. **These use cases were identified through extensive discussions with XNAT users and the broader imaging community as the emerging imaging scenarios most in need of informatics support**. To support these use cases we propose to complete the following aims:

***Aim 1. Execute the XNAT development roadmap*.** We will implement improvements to existing XNAT capabilities and introduce a number of new capabilities that enable our key uses cases. These capabilities include deeper integration with clinical information systems, scalability services including integration with the cloud, on-demand access to statistical and other computational platforms, a GUI-based form building interface, configurable data feeds, and quality assurance and communications tools. In addition, the core XNAT web application framework will be modernized, the user interface will be completely redesigned around customizable dashboards, and a “one-click” installer and an intelligent download application will be implemented. Many of these capabilities will be implemented using XNAT’s new plugin architecture and distributed on the newly developed XNAT Marketplace website, which together enable a rapid release schedule and “profiles” of related capabilities configured as a single package.

***Aim 2. Develop and evaluate custom XNAT profiles for each key use case*.** Our four key use cases each present unique requirements. We will capitalize on the platform’s support for profiles to develop XNAT configurations for each use case. Each use case profile will include a suite of plugins and custom dashboards specifically tailored to the unique aspects of the use case. Each profile will be carefully tested and validated in real world scientific projects. The institutional repository profile will focus on leveraging the shared information and computing resources within the institution. The clinical research profile will focus on a bidirectional workflow between clinical system and XNAT’s computational services. The multi-center trial profile will focus on focus on data integrity and quality assurance. Finally, the data sharing profile will focus on scalable dissemination of research data.

Aim 3. Provide support to the XNAT community. Active support of the XNAT community will focus on four categories that cover the breadth of XNAT’s constituents: end users, system administrators, developers, data miners, and principal investigators. Specific support and outreach initiatives will be developed to target each user group. End user support will by primarily through help sections built directly into the application as well as online tutorials and instructional videos. System administrators will be supported through an active online discussion groups, annual workshops, and online documentation. Developers will be supported through the online discussion group, in person “hackathons”, and a variety of developer tools and documentation. Principal investigators will be informed of XNAT capabilities and opportunities through peer-reviewed papers, conference presentations, and an active outreach programs.

These aims represent a significant expansion of the XNAT program, including in the breadth of covered use cases and in the extent of outreach to the imaging research community, an expansion that we believe is warranted by the quality of the platform and the need in the community.

SIGNIFICANCE

Why imaging? The role of biomedical imaging in basic and translational research continues to grow at a high rate (**Figure 1**) [1,2]. In his speech at the "NIBIB 5th Anniversary Symposium" in 2007, Dr. Elias Zerhouni, then director of the National Institutes of Health (NIH), identified biomedical imaging as the primary tool by which "biological data of the future" will be obtained [3]. Five years later, that future is now. In fiscal year 2012, the NIH is funding 1,939 imaging-related projects, encompassing virtually every institute and totaling over $730 million in total costs (based on NIH RePORT data extracted using the project terms "magnetic resonance imaging", "positron emission tomography", and "computed tomography").

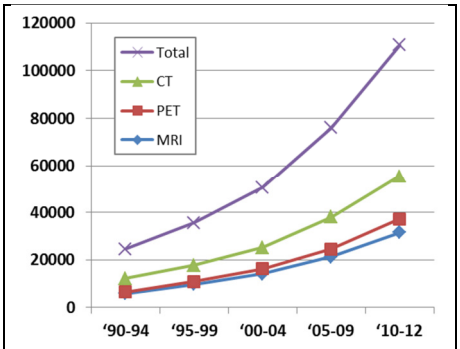

Figure 1. Growth in biomedical imaging research. Average annual publications per year, based on a Pubmed search using the modality abbreviation as the search term. For comparison, the keyword "bone" barely doubled from 1990 to 2012.

Why imaging informatics? While imaging is not unique in its importance in research, a number of attributes of imaging research pose unique data management challenges:

✓***Imaging data are very large.*** In a typical study, individual subject data range up to several gigabytes, and the study may include hundreds or thousands of subjects. Computational methods can expand the overall data set by an order of magnitude. Informatics tools can enable efficient storage and access to these data.

✓***Imaging requires extensive processing and analysis.*** Analytic procedures must be carefully managed to ensure accuracy, reliability, and integrity. Quantitative measures, which are increasingly the standard in imaging, carry the additional complication of relying on algorithmically sophisticated and computationally intensive processing streams. Informatics tools can provide infrastructure for executing systematic, verifiable post-processing and analysis procedures.

✓***Imaging studies are highly cross-disciplinary.*** Imaging is often one of many phenotypic assays that are being used within a research protocol. Correlations between imaging and these other measures may be calculated based on summary image measures or increasingly in a voxel-wise manner across an entire image. Informatics tools can enable coherent aggregation and statistical analysis across the full range of a study's data.

✓***Imaging studies rely on complex instrumentation.*** Whether in the context of a multi-center trial or in a single institution, images within a one study are often obtained on multiple different scanners using similar but not necessarily identical acquisition sequences. Even on a single scanner, sequence versions and underlying hardware change over time. Informatics tools can play a role in aggregating data from multiple locations and ensuring that data are acquired uniformly across sites and over time.

Given these unique challenges, especially as they build on the broader data management challenges presented by modern biomedical science (patient privacy, security, translational research, data sharing), *imaging informatics tools are vital to supporting imaging research*.

Why XNAT? XNAT provides a broad set of capabilities to address the issues noted above. While other systems provide similar features as XNAT's [4-8], XNAT is unique in number of ways:

✓***XNAT is part of the international informatics backbone.*** XNAT is a core component of the national and international programs that are driving informatics innovation, including the Biomedical Informatics Research Network (BIRN), the National Alliance for Medical Image Computing (NAMIC), Informatics for Integrating Bench to Bedside (I2B2), The Cancer Image Archive (TCIA), and the Human Connectome Project.

✓***XNAT is open source.*** XNAT has been open source since its

What exactly is XNAT?

XNAT is a web-based software platform designed to facilitate common management and productivity tasks for *in vivo* imaging and associated data. It consists of an image repository to store raw and post-processed images, a database to store metadata and non-imaging measures, and user interface tools for accessing, querying, visualizing, and exploring data. XNAT supports all common imaging methods (e.g. MRI, CT, PET), and its data model can be extended to capture virtually any related metadata (e.g. demographics, genetics). XNAT includes a DICOM workflow to enable exams to be sent directly from scanners, PACS, and other DICOM devices. XNAT's web application provides a number of quality control and productivity features, including data entry forms, searching, reports of experimental data, upload/download tools, access to standard laboratory processing pipelines, and an online image viewer. A fine- grained access control system ensures that users are restricted to accessing only authorized data. XNAT also includes a web services API for programmatic access and an open plugin architecture for extending XNAT's core capabilities. A number of tools and plugins are available on the XNAT marketplace website.

inception and has an active community of developers contributing to the core code base, building XNAT extensions, and creating new XNAT tools.

✓***XNAT is a platform.*** XNAT employs open interfaces, open packaging mechanisms, and open distribution channels to enable developers to hook into the system's core capabilities. A core aim of the current funding period is to implement "an XNAT-supported imaging informatics ecosystem… that will speed scientific discovery and bench-to-bedside translation". With the recent development of the XNAT REST API and the XNAT plugins infrastructure, just such an ecosystem has emerged.

✓***XNAT is used across the biomedical spectrum.*** XNAT has emerged as the most widely used informatics system in the field. To provide a sense of the breadth of programs utilizing XNAT, some example XNAT instances are listed in **Table 1**. *Letters of support detailing these projects and three dozen other XNAT-dependent research programs have been provided with this application.*

Domain	Program/Institution	Investigator(s)
Cancer	QIN – MGH	Rosen
Cancer	QIN – BWH	Fennessy
Cancer	QIN – Iowa	Buatti
Cancer	The Cancer Image Archive	Freymann
Cancer	Institute of Cancer Research	Doran
Neuroscience	Human Connectome Project	Van Essen
Neuroscience	Harvard University	Buckner
Neuroscience	NITRC	Preuss
Ophthalmology	University of Iowa	Abramoff
Cardiology	Cardiovascular Research Grid	Winslow
Cardiology	Wake Forest University	Carr
Multidisciplinary	Johns Hopkins University	Miller
Multidisciplinary	CTSA, Harvard Medical School	Rosen
Multidisciplinary	CTSA, Washington Univ.	McKinstry

Table 1. Example XNAT deployments. These examples and those described in over three dozen additional letters of support illustrate the breadth of science supported by XNAT. Many additional XNAT installations are in operation at institutions around the world.

Why this grant? Given XNAT's unique position as the premier open source imaging informatics platform, we believe that the renewal of this grant is essential to ensure that XNAT continues to provide the scientific community with exceptional and relevant capabilities. In addition, we propose to significantly expand XNAT's feature set to enable new and what we believe to be truly transformational modes of scientific research that break down walls between scientists, institutions, and the clinic. While these new initiatives are ambitious, they are backed by a highly experienced principal investigator and development team, and they are grounded by the ongoing guidance provided by close interaction with real world scientific programs.

INNOVATION

Some aspects of this proposal are deliberately *not* innovative. Indeed, the XNAT technology stack is based on widely used open source tools that are stable and well supported. Our testing, documentation, and support methods follow widely used industry best practices. These deliberately non-innovative tactics ensure that XNAT is a stable, rock solid application. So while not innovative in themselves, our tactics yield an extremely powerful platform that enables highly innovative and high impact science. Furthermore, there are many aspects of this proposal are highly innovative including technical innovations to integrate with clinical information systems and to dynamically scale systems to support large data sharing projects. The proposed Data Feed Engine and Computational Engine will enable a new and exceptionally powerful breed of analytic approaches to mine image repositories.

APPROACH

Research team. The research team is led by Dr. Daniel Marcus, who has directed the XNAT program since its inception. As the PI of this and a number of other service and scientific programs, Dr. Marcus has considerable experience leading informatics and imaging science programs. Dr. Lawrence Tarbox and Mr. Steve are leading experts in medical imaging formats and data exchange protocols. Mr. Richard Herrick leads a team of highly skilled software developers who have been contributing to the XNAT codebase for an average of 2.5 years. The team is augmented by Timothy Olsen at Deck Five Consulting, who was the lead XNAT architect from 2004 to early 2012, and Integrity STL, an award-winning digital strategy and design firm. More detailed credentials for each team member are provided in the Budget Justification and biosketch personal statements.

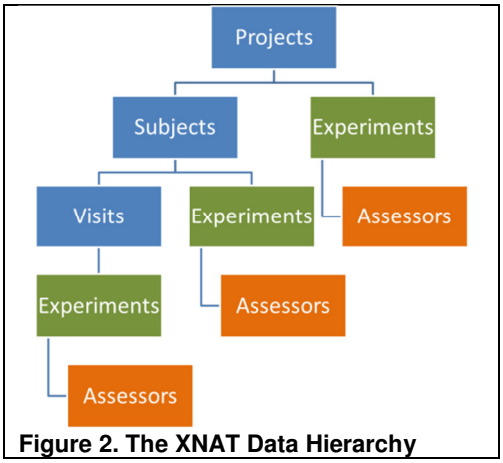

Figure 2. The XNAT Data Hierarchy

Recent progress. "The XNAT Imaging Informatics Platform" grant was awarded in August, 2009. Over the next 3 years, XNAT versions 1.4, 1.5, and 1.6 have been released. With each successive release,

XNAT's core capabilities have been expanded, including its support for data structures, navigation, security, and workflow. Some of the highest impact capabilities include:

- *Hierarchical data model*. XNAT's hierarchical data model (**Figure 2**) matches the organizational principles commonly used in human imaging studies. This hierarchical data model is leveraged throughout XNAT, including in its core navigation interface and in the organization of its API. At the root, the Project entity drives much of XNAT's organization and operation, including the security infrastructure.
- *Web-based user interface (UI).* The XNAT UI (**Figure 3**) has undergone signficant revision in recent versions to improve usability and speed. The UI includes navigation paths based on the XNAT data hierarchy and by specific data types, detailed reports of individual experiment data, and basic and advanced search interfaces.

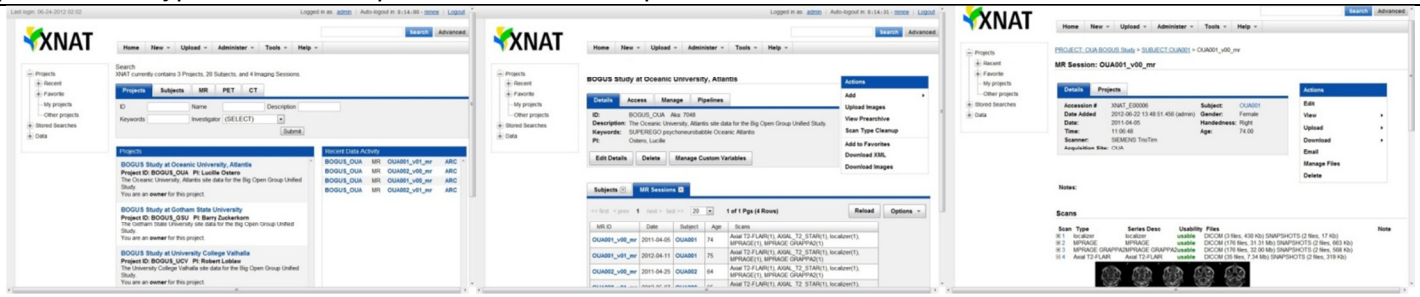

Figure 3. The XNAT web interface. Example pages include (from left) home page, an individual project page, and a detailed experiment page (MR Session in this case). Breadcrumbs provide hierarchical navigation paths while the left-hand menu links to views by data type. Additional pages for data entry, search, data tables, administration, and other features are not shown due to space limitations. Note: Larger versions of all figures are provided in the appendix.

- *Study management*. XNAT includes a variety of features to support the execution of research studies in a variety of contexts, including prospective studies, retrospective studies, longitudinal studies, and multi-center studies. Central amongst these, XNAT's Protocol Tracker provides investigators with a mechanism to define their experimental protocols and from these definitions track their study data. XNAT uses the protocol definitions to tailor data entry workflows, organize reports, and notify users of protocol violations.
- *DICOM Workflow*. With XNAT's DICOM C-STORE service, data can be directly archived from all DICOM modalities. The C-STORE service directs incoming data to appropriate projects based on values extracted from the DICOM metadata. The mapping logic for routing data to projects is highly configurable. During the archive process, site-wide and project-specific anonymization profiles are applied to the data, a configurable set of DICOM tags are extracted into the XNAT database, and a validation service is executed to evaluate whether incoming data comply with the project's protocol. XNAT also includes a DICOM review screen for manually reviewing DICOM tags for discrepancies and patient identifying information. XNAT's DICOM workflow can also be paired with DICOM Browser [9] to achieve more complex workflows.

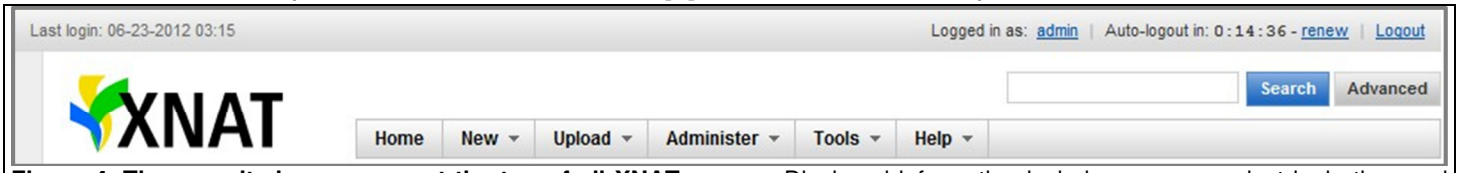

Figure 4. The security bar appears at the top of all XNAT pages. Displayed information includes username, last login time and location, current session time remaining (which renews with each page submit), and security warnings.

- *Security*. XNAT's security infrastructure was dramatically revised in XNAT 1.6 both to provide more security features and to prevent various security threats. The underlying user authorization implementation was ported to Spring Security, which supports multiple authentication protocols, including XNAT's local database, LDAP, Active Directory, and OpenID. New features include configurable password complexity and expiration, multiple login detection, and a security bar UI component (**Figure 4**). Additional security capabilities were added to address the OWASP Top Ten threats and vulnerabilities, including cross-site scripting, SQL injection, and code injection [10]. All XNAT development is now tested using one or more web vulnerability evaluation tools (e.g. BurpSuite Pro). Prior to full release, XNAT 1.6 passed an independent security audit conducted by the University of Texas Chief Security Office (see letter provided by R. Poldrack).
- *Quality assurance*: XNAT includes a variety of quality assurance features, including a "prearchive" for quarantining incoming data, forms and data types for manual review of images, and pipeline support for automatic image quality review. XNAT 1.6 supports a full audit trail on all transactions, from data creation and manipulation to user configuration and data access attempts.
- *RESTful API*. A web services programming interface (API) based on the REST model [11] was released in XNAT 1.4 and expanded with each subsequent release [12]. The API provides comprehensive access to

•XNAT services, including those for storing and retrieving data, querying the database, performing administrative tasks, and running pipelines. The API is used extensively throughout XNAT itself and also by an array of external applications that integrate with XNAT.

•*XNAT Plugins and Marketplace.* XNAT 1.6 includes a plugin architecture that allows developers to bundle all of the required components for a new capability – data model, web pages, java classes, API extensions -- into a single package that can be automatically imported and deployed by XNAT. This plugin architecture is a major advance, as it enables new XNAT features to be created and distributed independent of the typical release schedule. We developed the XNAT Marketplace website (http://marketplace.xnat.org) to serve as a formal channel for distributing these plugins as well as XNAT-enabled tools and scripts. The site follows design patterns pioneered by the Mozilla Foundation for organizing, tagging, and featuring content (**Figure 5**). We anticipate that the XNAT Marketplace will invigorate the XNAT developer community and will significantly expand the overall XNAT ecosystem.

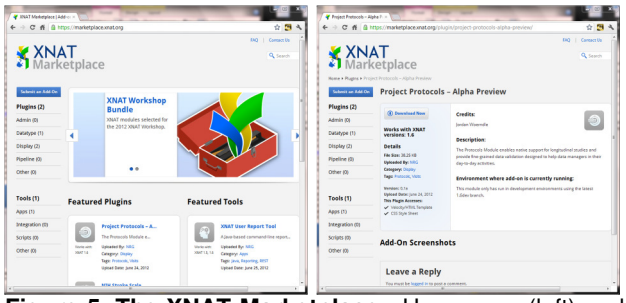

Figure 5. The XNAT Marketplace. Homepage (left) and an example plugin (right). Currently, only a handful of components have been added prior to its scheduled public release in August, 2012. Many more plugins will be contributed in the coming months.

SPECIFIC AIM 1. Execute the XNAT development roadmap

The major deliverables on the XNAT roadmap are described below, organized by functional category. For each deliverable, a number is provided to indicate the XNAT version (1.7, 2.0, 2.1) or plugin package (#1-#6) in which the deliverable will be included. A schedule for the versioned releases and plugin packages is provided in the Timeline section.

Core components. Ongoing development of the XNAT core will focus on modernizing and streamlining a variety of internals components and incorporating additional services.

•*Spring Framework and Hibernate (2.0):* The XNAT web application currently relies on the Apache Turbine web application framework, which has proven to be remarkably stable but is no longer actively supported. We will migrate the web application framework to the Spring Framework, which has emerged in recent years as the most widely used and feature-rich Java framework. In addition, the component that communicates between the web application and the database, which currently utilizes custom XNAT code, will be replaced with Hibernate, an industry standard for object-relational mapping. Moving to Spring Framework and Hibernate comes with many benefits, including stability, an expanded feature set, and active support mechanisms. In addition, because these tools are so widely used, the learning curve for both internal and external XNAT developers will be greatly reduced.

•*Dynamic Plugins (1.7)*. While XNAT's current plugin architecture is a major advance, it has some limitations. In particular, the plugin packages must be manually downloaded and placed into a plugins directory. The web application must then be rebuilt before the plugin is active. The proposed dynamic version of the plugin system will allow users to find plugins on XNAT Marketplace and click an "Install" link to instantaneously deploy and activate the plugin.

•*Rules Engine (2.1).* The XNAT Rules Engine will enable rules related to XNAT-hosted data to be defined and to trigger various actions. For example, a user could define a rule to disallow enrollment of control subjects when the desired enrollment target has been met. The Rules Engine will be used widely within XNAT to trigger actions and a user interface will allow users to set their own rules. The system will build on Drools, a widely used business logic platform that provides a declarative language and engine for defining and executing rules.

•*Notification Engine (2.0).* A Notification Engine will be implemented to serve as a messaging hub for XNAT activities. For example, users could request notifications whenever an image processing pipeline they have launched completes. Notifications will be tied closely with the Rules Engine, so that users can request to be notified whenever a particular rule submits a message to the Notification Engine. A central message center will be built into the XNAT UI, and the engine will support a variety of additional notification methods, including RSS feeds, email, text, and tweets. The notification engine will extend XNAT's current synchronous notification tool to include asynchronous messages and Java Messaging Service queues.

•*Data Feed Engine (1.7)*. A Data Feed Engine will be developed to feed external data visualization and analysis tools with XNAT-hosted data. The engine will extend XNAT's existing API to serve data in formats preferred by these tools. The Data Feed Engine will be accessible using the "format" parameter that is already

part of the API. For example, to retrieve a Google Charts-formatted list of subjects with age and gender, the API call would look like: https://central.xnat.org/data/subjects?columns=ID,gender,age&format=google_chart. XNAT will come with a library of data feeds for widely used tools, including Excel, SAS, Google Charts, D3, and Highcharts. To enable the service to support additional formats, we will implement a template parser that dynamically generates data feeds based on a template provided with the HTTP request. For more complex formats, custom java classes that produce particular data formats can be developed and deployed as XNAT plugins. Developers will be able to use the Data Feed Engine to generate interactive charts on XNAT pages and to embed XNAT content in external applications. A prototype of this service has been implemented to feed XNAT data into a D3 chart (**Figure 6**).

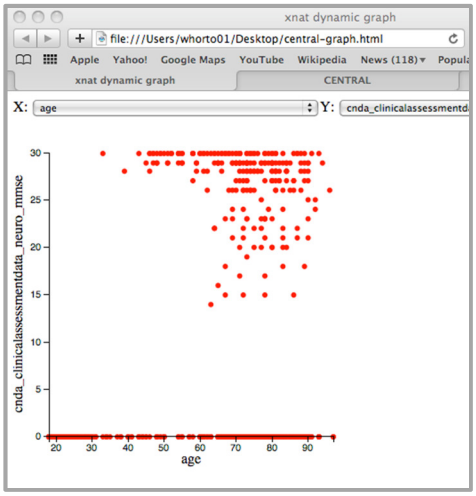

Figure 6. Data Feed Prototype dynamically populates a D3 chart with age and MMSE scores from an XNAT data feed based on a single API call to XNAT Central and could be easily embedded within an XNAT site or in an external application. The code is at https://gist.github.com/3018714.

File management. The bulk of data managed by XNAT is in files on the file system. These files include the acquired data and the various derived data generated by image processing pipelines. Given the importance of these files, file management is an essential XNAT capability.

- *File Integrity Service (#1)*. Currently, XNAT provides a REST API resource for every file it manages. The File Integrity service will add a number of optional parameters to this resource, including an MD5 checksum, version information, access to historical versions, provenance information (using the W3C prov model [13]), and user-supplied tags. An automated format conversion parameter will also be added to retrieve files in dynamically generated alternative formats. For example, a DICOM file could be requested as a JPG image.
- *File Download Manager*. Because image data are so large, typical web-based download methods, such as zip-formatted archive files are grossly insufficient. XNAT currently includes a Java applet-based tool for managing downloads, but this too has limitations in its ability to preserve information between sessions. For example, if a download stream is broken midway, the applet must start the download over from the beginning. We will develop an installable Java-based download manager built on the Eclipse Rich Client Platform. The download manager will include a number of features for streamlining downloads, including file verification by MD5 checksum, start/stop/pause/restart capabilities, synchronization of local data to XNAT repositories, schedulable download times, and batch download of data by project and search result. The download manager will utilize the existing file API and the File Integrity Service described above.

User interface (2.0). The XNAT user interface will be rewritten to utilize current web technology and build on best practices for user experience. The redesign process was initiated in November 2011 through a series of user focus groups and surveys conducted by Integrity, an award winning digital strategy firm. As a result of this community input, in partnership with Integrity, we have produced a full functional requirements specification [14] and draft designs [15] (**Figure 7**). While the redesigned interface includes significant new look and feel, many of the existing design elements that users rated highly are retained, including the hierarchical navigation, filterable data tables, and individual data reports. Within the scope of this proposal, these designs will be fully implemented in XNAT. Several new features are worth noting. **Configurable Dashboards** enable system administrators and users to tailor pages to their specific workflows. For example, an administrator could build a dashboard for MRI technologists that includes quick links to upload tools and a review pane for recent studies sent from their scanners. The homepage dashboard can be configured by the administrator for specific users based on their assigned role in the system. A **Navigation Bar** at the top of every page provides access to key site elements while consuming minimal page space. **Task Ribbons** near the top of each page provide context-specific actions, similar in style to recent versions of Microsoft office. A **Data Table Wizard** provides users with an intuitive interface for generating complex tables of data. Finally, overall page layout was designed specifically to suit XNAT's plugin architecture. For example, each page is laid out in a three column design. Plugins can be designed to consume one, two, or three columns on the page. A page layout manager handles the specifics of positioning the plugin on the page.

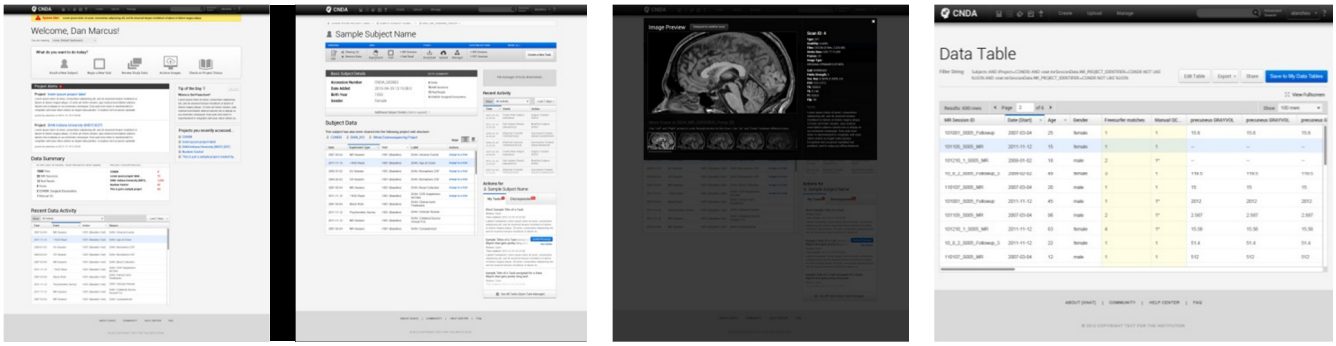

Figure 7. The XNAT 2.0 web interface design. These examples of (from left) the default homepage dashboard, subject report, image viewer, and data table illustrate the look and feel of the redesigned UI. Clickable mockups of the full site are available at http://integrity.protosharesites.com/i8529051560561775385 6/#, and a preview of page designs is available at http://integrity.prevuu.com/p/xnat.

In addition to implementing the redesigned web interface, we propose to develop three additional UI services:

- *Form Builder (#6).* A Form Builder service will be developed to enable users to create case report forms and other data entry forms through a GUI layout manager. The Form Builder will be designed to work with XNAT's current XML Schema-based data extensions model and with a dynamic extensions model for developing ad hoc data types. The Form Builder will allow users to select appropriate form elements (e.g. checkboxes, radio buttons, dropdown menus) and to position them on the screen. Similar capabilities have proven successful in a variety of web-based data entry systems (e.g. REDCap) and we expect will be suitable for designing the most forms required in typical studies.
- *Discrepancy Manager (#2).* Tracking and resolving data discrepancies (often referred to as "queries" in clinical trial parlance) is a core quality assurance task in clinical research. We will develop a Discrepancy Manager interface to assist coordinators in tracking and resolving discrepancies in uploaded and entered data. Discrepancies will be generated by the existing Protocol Tracker, by the proposed Rules Engine, and by coordinators screening data. The Discrepancy Manager will log all discrepancies, track the "conversation" between parties as the discrepancy is resolved, and provide various QA metrics (e.g. time to resolution, discrepancies/site).
- *Report Engine (#5).* We will develop a Report Engine to enable generation of custom reports using Jasper Reports, an open source report generation system that renders reports into many common formats including HTML, PDF, Excel, and Word. We will develop and distribute a library of preconfigured reports with XNAT, including NIH-formatted enrollment tables, key performance indicators, and quality assurance summaries. Site managers and users will be able to develop and upload custom reports using the iReport report designer.

Administration. Currently, XNAT deployment requires installation of prerequisite software (Tomcat, PostgreSQL), execution of a build script, and import of the built application into Tomcat. Ongoing maintenance requires modifying configuration files within the web application package and inspecting a byzantine set of log files. Because the difficulty of these tasks has proven to be a limiting factor in XNAT's use at many sites, we will invest considerable effort in streamlining them.

- *One Click Installer (2.0).* We will develop an installer package that removes all external dependencies. The installer will include an embedded web server (Jetty) and in-memory database system (H2). After running the installer, the site manager will navigate to the newly installed XNAT in his browser and complete the configuration process through a series of well-documented webpages. We expect that the embedded web server and database will be performant for XNAT systems with traffic on the order of 50-100 concurrent users. For larger systems, we will provide migration scripts to rollover to Tomcat and PostgreSQL to enable higher throughput.
- *Administration Center (#3).* We will develop an Administration Center on the web application that provides a unified, GUI-based interface to all administrative tasks. A configuration panel will enable editing of all XNAT settings, including DICOM configuration, site-wide anonimization profiles, backup procedures, and pipeline services. A monitoring panel will provide access to site usage and load with extensive analytics provided by Google Analytics. A diagnostics panel will provide full access to logging information and assist site administrators in resolving site issues.

Scalability. As imaging repositories expand to include more data and to serve more users, XNAT's ability to scale will become critical. We will take a multi-pronged approach to ensuring that XNAT is massively scalable.

•*Load balancing (1.7)*. XNAT currently runs within a single web server, which severely limits throughput. We will use Apache HTTP 2.2's mod_proxy and mod_proxy_balancer to enable load balancing across multiple backend web servers. In addition, the data caching mechanism methods within XNAT will be refactored to use Ehcache, an open source API that will enable a user's session state to be accessed by multiple independent XNAT servers.

•*Distributed file systems (2.1)*. We will implement support for distributed file systems, including Amazon S3, Google Cloud Storage, and iRODS. XNAT's existing file management infrastructure, which utilizes "catalogs" of URL-based file references, is quite amenable to supporting these systems. These URLs can point to either local files or remote web resources. Currently, XNAT downloads remote files and then passes the file on to the user over the REST API. The more efficient approach would be for the user to download the remote file directly. To enable this, we will modify the REST API to allow users to request the resource URL instead of the actual resource. An Oauth token will be provided with the URL that grants the user access to the file.

•*Data Mirroring (2.0)*. This service will enable copies of data to be synchronized to one or more mirror locations and made accessible via a single XNAT web interface. The mirror locations will each include an independent XNAT system which will be used to host data but will not be accessed directly by users. Synchronization will be managed by the primary location, which will be triggered by the Rules Engine to initiate copies of data to each mirror site whenever data are added or modified. Data copies will be executed using the REST API. File time stamps and MD5 checksums will be used to determine which files need updating. Mirroring will be configurable on a per project basis. A UI will be developed on the web application to enable project managers to select mirror sites, configure mirror settings, and monitor mirroring status.

•*Cloud Backup Service (#3).* This service will securely back up data to the cloud. Backups will be configurable on a per-project basis, allowing project owners to select cloud vendors and provide account information through the project management UI. We have implemented a prototype of this service that encrypts and synchronizes project data to a ZFS-based storage system on an Amazon EC2 instance with dynamically expandable block storage. Data synchronization is enabled via scheduled execution of the ZFS Send command which efficiently transmits any changes since the last system snapshot.

Computational Engine (2.1). Computation is essential to XNAT, as it enables the quantitative image processing and analysis that are at the heart of our key use cases. We propose to develop a new computational engine in XNAT that enables managed, scalable execution of computational routines on any analysis package, including widely used platforms (and locally developed tools. The engine will be composed of several components. First, the **Toole Registry** will manage the set of specific tools that are available within a particular XNAT instance (e.g the R statistics platform). A UI for the Tools Registry will be provided for site administrators to add and remove packages from their instance's registry. Second, the **Routines Registry** will manage the set of specific scripts that are available to users of the site (e.g. an R script to calculate mean lesion volumes in a patient group). The Routines Registry will include both a global library, managed by the site administrator, and personal libraries managed by individual users. Each routine in the registry will include several metadata attributes: name, description, estimated execution time, associated package, and an XNAT Data Feed format. The Data Feed format is the essential component that ties generic computational routines (the R script) with XNAT-hosted data (the list of patients). Third, the **Compute Manager** will serve as the central bus for managing requests for compute services. The Compute Manager will include an API for posting service requests (e.g. run the mean lesion calculation script on this set of patients), monitoring their execution, and retrieving cached results. Finally, the Compute Manager will manage a **Compute Pool**, a sandboxed set of virtual machines on which the service requests will be executed. When a service request is received, the Compute Manager will execute the job on a Compute Pool node. Events will be posted to the Notification Engine when jobs are launched and completed.

DICOM. We will enhance XNAT's DICOM services to support a much broader workflow.

•*Modality worklist (#5)*. Modality worklists will be built into XNAT to enable users to push acquisition protocols to a scanner. Protocols will be managed on a per projects basis, allowing help investigators ensure that a consistent protocol is used within their studies.

•*Push/pull (#1)*. XNAT will also act as a Service Class User (SCU), which will allow it to push and pull data from modalities, PACS, and other DICOM devices. This capability enables a number of use cases, including retrieving studies from a clinical PACS, applying an anoymization profile, and assigning the studies to a project in XNAT. Processing pipelines can then be run on these studies and the resulting secondary images can be pushed back to the PACS to enable clinical translational studies.

•*Format conversion (#4)*. Research post-processing applications often output non-DICOM images. The format conversion service will enable these images to be converted into either DICOM secondary image captures, DICOM-RT objects, or DICOM structured reports, as appropriate within the particular image context. Combined with the DICOM push service described above, this will be a powerful way to send post-processed data to clinical PACS and workstations.

Integration Activities. XNAT is most powerful when joined with complimentary applications such as visualization tools, data analysis suites, data mining tools, and related databases. The XNAT API and plugin architecture have made this type of integration very straightforward. Some examples of XNAT extensions and XNAT-enabled applications are listed in **Table 2** and detailed in letters provided by their developers and in recent publications [16-21]. We propose to undertake a number of initiatives to further these integration activities.

•*Database plugins (#3)*. We will develop XNAT plugins for I2B2 and REDCap, two widely used database systems for mining medical records and capturing form-based data, respectively. Significant progress towards I2B2 integration has already been made via the mi2b2 project (see letter from S. Murphy). Here, we will enable access to I2B2-hosted clinical data for patients in XNAT. For both I2B2 and REDCap, we will use the existing APIs provided by the platforms and will enable single sign-on capabilities between XNAT and the external system.

•*Visualization modules*. We will develop modules in popular visualization tools to enable input/output operations between XNAT and the tools. This type of interaction has already been implemented in several tools, including 3D Slicer, CAWorks, and Connectome Viewer. Additional tools that will be enabled are ImageJ, XIP, MIPAV, Osirix and Mango. In implementing interfaces with these tools, we will develop a library of code and common UI elements (e.g. tree views, file selector dialogs) that will reduce the effort required to integrate each additional tool.

• *AIM import service (#1)*. The Annotation and Image Markup (AIM) language for representing image annotations, including regions of interest and lesion measurements, is emerging as a standard. Working closely with the AIM developers (see letter from E. Siegel), we will implement an AIM import service. The service will extract annotations from imported documents and associated AIM templates and store them to the XNAT database, which will enable searching on detailed annotation data.

•*Content Management System (CMS) plugins (#6)*. CMS and portal platforms provide rapid development environments for building sophisticated, data-rich websites. We will develop plugins in Plone, Drupal, and Liferay, three such systems that are in use at XNAT sites. The Cardiovascular Research Grid has already developed an initial version of a Liferay XNAT plugin that we will build on (see letter from R. Winslow).

Tool/Module	Project Lead
PyXNAT	Schwartz
3D Slicer	Kikinis
XNAT DataChooser	Doran
XNAT Viewer	Gutman
LONI Pipeline	Van Horn
Grid Wizard Engine	Grethe
Connectome Workbench	Van Essen
Mi2b2	Murphy
INCF One-Click Uploader	Hill
BOLD QC	Buckner
Mango	Lancaster
BIRN Mediator	Kesselman
Connectome Viewer	Gerhard
CAWorks	Miller

Table 2. The XNAT Ecosystem. Examples of XNAT-enabled tools developed by the community.

Specific Aim 2. Develop and evaluate custom XNAT profiles for each key use case.

The capabilities described in Aim 1 serve as the foundation for enabling our four key use cases. In Aim 2, we will compose these capabilities into tailored profiles specifically designed to support each of our use cases. For each use case, we have selected one or more real world programs that will be used to model and test these configurations. We will work closely with the scientific and technical staff from each of these projects from inception through ongoing production. In a number of cases, XNAT is already in use by the project, and the project staff has already contributed valuable feedback to the XNAT team. A letter of support has been provided by each of the project PIs, indicating his enthusiasm to engage with us. This approach of tying XNAT development to real world science has been used since XNAT's earliest incarnation in the Washington University Alzheimer's Center and we believe has contributed significantly to XNAT's success. Thus, this aim achieves two primary goals: 1) to deliver useful use-case specific XNAT configurations and to provide ongoing feedback to the XNAT development efforts.

Below, we describe the high level requirements and current capabilities that relate to these requirements for each use case. We then propose specific configuration profiles to fully address the use case and detail how these profiles will be evaluated in real world projects.

Institutional Repositories

Overview. Institutional repositories provide a single centralized resource to host imaging and related data from many investigators. Central repositories have a number of benefits include efficiencies of scale and efficient integration with data resources like PACS and databases in related realms. A central repository also encourages reuse of research methods, including image acquisition protocols and clinical forms.

Current Capabilities. Amongst XNAT's capabilities for supporting institutional repositories, the project-based organization and access control system is most critical, as it mitigates the common concern investigators express over "losing control" of their data. XNAT's DICOM C-Store workflow provides a reliable and automated mechanism for capturing and routing data from an institution's imaging facilities, and XNAT's flexible file management and pipeline services enable it to utilize institutional network storage and computing systems.

Proposed Profile. We will configure an XNAT profile for institutional repositories that focuses on the availability of shared resources within the institution. The DICOM worklist service will enable investigators to select from common acquisition protocols that are used at the institution, assign the protocol to a particular project, and then push the protocol to the scanner when a subject is scanned. The DICOM validation service will be configured to verify that the correct protocol was indeed executed. The XNAT authentication service will be configured to utilize institutional identify management systems, typically through LDAP or Shibboleth, which will allow users to login using their standard username and password. Custom dashboards will be developed specifically for investigators who manage multiple studies within the system and for imaging technologists to perform quality assurance reviews of the data from the scanners they manage. Finally, a library of commonly used clinical and behavioral forms will be included to encourage use of standard data elements across studies and between investigators. The Form Builder service will be provided for creating forms that are not available in the library.

Evaluation Projects. The Clinical and Translational Science Award (CTSA) programs at University of Iowa and Washington University will serve as our models for institutional repositories (see letters from B. Knosp and B. McKinstry). Both programs currently operate large scale XNAT deployments that are accessible to their investigators. At Washington University, the deployment will be linked to the University's WUSTLKey Shibboleth authentication system and to caBIG and REDCap-based clinical and behavioral databases. A standard brain imaging MRI protocol that's widely used by Wash U investigators will be built as a worklist and paired with an associated validation profile. Similar tasks will be completed at U. of Iowa, using their HawkID authentication system and REDCap databases.

Clinical Imaging Research

Overview. Clinical imaging research includes both *retrospective studies*, which rely on batches of historic patient scans meeting some diagnostic criteria, and *prospective studies*, which rely on individual patients scans to be imported and often processed in real-time. From an imaging informatics perspective, support for clinical research requires integration with clinical devices and information systems, careful compliance with regulatory requirements; and agility in moving between clinical and research data formats and protocols.

Institution	Site Lead
Duke University	Daniel Sullivan, Vice Chair of Radiology
Emory University	Allan Levey, Chair of Neurology
Johns Hopkins University	Michael Miller, Director, Center for Imaging Science
Massachusetts General Hospital	Bruce Rosen, Director, Martinos Center for Biomed. Imaging
University of Washington	T. Grabowski, Director, Integrated Brain Imaging Center
Washington University	R. Gilbert Jost, Chair of Radiology

Table 3. Clinical Research Early Adopters Program.

Current Capabilities. XNAT's existing DICOM workflow, pipeline service, and a variety of standard clinical forms (e.g. Radiological reads, NIH Stroke scale) provide support for clinical imaging research. XNAT's overall security infrastructure is well-suited to clinical research.

Proposed Profile. We will configure an XNAT profile for clinical research that focuses on a secure, bidirectional workflow between XNAT and clinical systems. The profile will include several methods for obtaining clinical scans, including direct sends from scanners and query/retrieve from PACS. XNAT's I2B2 plugin will provide a graceful interface for mining electronic medical records to discover patients for retrospective studies. The plugin will feed query results directly to XNAT's DICOM query/retrieve service. Using XNAT's anonymization service, patient information in the image files will be encrypted in the XNAT archive and pipelines. XNAT's Computational Engine will be used to enable rapid and reliable image processing and analysis which will be verifiable using the File Integrity Service. The DICOM Formatting service will be configured to convert files generated by pipelines from research formats to the appropriate DICOM objects (secondary captures, structured reports, RT objects). Encrypted patient information will be restored as derived images are pushed back to the PACS through XNAT's DICOM Push service. Finally, custom dashboards will be created to enable

investigators and clinicians to track patient information in XNAT in a style that is similar to existing web-based patient records.

Evaluation Projects. We have formed a Clinical Research Early Adopters Program and invited an initial group of institutions to participate (**Table 3**). Dr. Marcus and the XNAT outreach coordinator will make an onsite visit to each institution to identify specific clinical research workflows that can be enabled at the institution. Each site will be provided with an XNAT server, preconfigured with the clinical profile, and the outreach coordinator will work with site staff to ensure that the system is deployed and fully operational. At each site, the system will be integrated with a research patient database (typically I2B2) and at least one translational image processing pipeline will be deployed. Through continued outreach, including teleconferences and additional onsite visits, we will monitor the system's use and provide ongoing guidance to site investigators and technical staff. Given the importance of translational imaging research, the technical challenges of integrating with clinical information systems, and the fact that we are entering largely uncharted terrain, we believe that this the enhanced level of support and outreach provided through the Early Adopters Program is well justified.

Multi-center trials

Overview. Due to the increasing prevalence of imaging-based biomarkers, multi-center imaging-based clinical trials are becoming increasingly commonplace. Multi-center trials are complicated operations and present technical and scientific challenges on a number of fronts: deploying functionally equivalent scan sequences across heterogeneous devices; defining and enforcing common procedures for subject recruitment and data acquisition; defining and enforcing policies for scanner maintenance and upgrades; securely uploading data from study sites; securely distributing data to investigators; and executing quality control procedures at short latency.

Current Capabilities. XNAT is currently used to manage a number of multi-site trials. XNAT's project architecture enables management of individual site data, including site-specific anonymization, DICOM validation, and processing pipelines. The Upload Applet enables streamlined uploads from remote sites. The protocol service is used to ensure that data are acquired on time and according to protocol. Image review forms are used to document quality control processes. XNAT's audit trail system records all transactions on data and files.

Proposed Profile. The multi-center trial profile will focus on data integrity, quality assurance, and study monitoring. In addition to the current capabilities, the Rules Engine, Discrepancy Manager, and Notification service will be configured to assist in identifying protocol discrepancies and tracking their resolution. The File Integrity service will be configured to record file provenance records. The Report service will be used to generate summary and quality assurance reports for study investigators. Finally, the Data Feed Engine will be configured to output study data sets according to formats provided by study statistical analysis plans. Two custom dashboards will be developed, one for study site coordinators that provides streamlined access to data entry and image upload forms and one for central coordinators to track data, notifications, and queries across all sites.

Evaluation Projects. PREDICT-HD and GENFI are ongoing multi-center trials that use XNAT and are enthusiastic to serve as evaluation projects (see letters from H. Johnson and S. Ourselin for details). We will assist both sites in upgrading their XNAT system to use the multi-center trial profile. Our evaluation process will include a number of metrics for assessing the impact on project workflow, including error rates, protocol compliance, and time to query resolution. We will also work closely with technical staff to configure the custom dashboards for optimal data entry and study tracking.

Data sharing

Overview. Data sharing entails an investigator distributing his data, either openly, semi-openly, or in closed collaborations. Large NIH studies are required to share data and many smaller projects have realized the benefits of sharing [22]. With the increasing prevalence of sharing, XNAT is being used more frequently in this context. However, image data sharing presents challenges both in data intake (anonymization, provenance, file organization) and distribution. Data distribution is particularly difficult due the large bandwidth requirements for moving high volume imaging studies.

Current Capabilities. XNAT's access control system allows investigators to make their data openly accessible to users of their XNAT instance, accessible by request, or completely closed. Its support for anonymization and DICOM metadata review help ensure subject privacy and compliance with HIPAA regulations. In XNAT 1.6, a new service was introduced to enable investigators to harmonize their scan labeling scheme with commonly used terms. Finally, its extensible data model allows investigators to share a variety of non-imaging data and derived image data with their imaging studies.

Proposed Profile. The data sharing profile will focus on scalable dissemination of research data following best practices for documentation, data provenance, and use of common data elements. The profile will include the proposed scalability capabilities, including load balancing, distributed file systems, and data mirroring. In addition, the Computational Engine will be configured to support dynamic analysis of the shared data. The file integrity services will be configured to capture data provenance records. Custom dashboards will be used to enable investigators to add extensive documentation related to their projects and the analytics service will enable tracking of project activity. The Notification and Rules Engines will be configured to enable users to subscribe to receive a notification when new data meeting desired criteria are shared. Finally, the OpenID authorization method will be implemented to allow users to register to data sharing site using preferred logins (e.g. Google, Facebook). While not directly part of the profile, the Download Manager application will be a core part of the overall data sharing use case. Additionally, the proposed integration activities will enable access to shared data through the various visualization tools.

Evaluation Projects. The Cancer Image Archive (TCIA) [23] and the Human Connectome Project (HCP) [24] will serve as our model projects (see letters from F. Prior and D. Van Essen). Each project provides unique data sharing attributes. TCIA includes ongoing contributions of data sets, typically including patient scans that require careful scrutiny. TCIA will make heavy use of the AIM import plugin and a variety of integrated visualization platforms, including 3D Slicer, Osiris, and XIP. The Human Connectome Project involves a single data set of massive proportions (> 1 petabyte) that will be distributed in quarterly releases as it is acquired. The HCP will use the data mirroring capability to dynamically expand its capacity when load increases around these quarterly releases [25]. The base systems at Washington University will be expanded to include additional geographically distributed systems at the University of Minnesota and Oxford University.

Distributing and Maintaining Use Case Profiles

Each specific profiles described above will be assembled as a set of bundled plugins and distributed on XNAT Marketplace. XNAT Marketplace was specifically designed to support such multi-plugin bundles. This approach will allow system managers to quickly review the included plugins and to add or remove plugins to customize the profile for their exact requirements. Individual XNAT deployments will also be able to support more than one installed profile. We anticipate that institutional repositories, for example, may also serve as clinical research hubs.

Additional use cases

Smaller single-investigator laboratories are seemingly ignored in the use cases covered here. However, we believe that these laboratories will be well served by the streamlined installation procedure and the web-based Administration Center that we have proposed. Several letters from directors of small labs have been provided (R. Aguilar and C. Butson are excellent examples), attesting to the benefits XNAT provides to their research and their enthusiasm for the proposed project.

Several additional uses are specifically *not* covered in the scope of this grant, including animal imaging, electrophysiological data, pathology, and microscopy. With some modification, XNAT is quite suitable for supporting these domains, and we look forward to identifying opportunities to support them.

SPECIFIC AIM 3. Support the XNAT community.

The XNAT team actively engages with the XNAT community, including developers, site managers, investigators, and end users, to enhance their experience and to gather feedback for our continuous improvement processes. Our support methods currently include web-based documentation, an active web-based discussion forum and in-person workshops. In keeping with the expanded scope of the XNAT platform, we propose to expand these practices and institute several additional support initiatives.

Documentation. Current XNAT documentation tends to focus on installation and management of XNAT and lacks adequate content for end users. *We will develop new documentation specifically aimed at end users*. This documentation will be hosted on the xnat.org domain, so that it can be easily updated, but will be embedded in the XNAT web application itself so that it is presented in a relevant context. In many cases the documentation will include brief instructional videos, typically less than 1 minute in length.

Workshops and Hackathons. The recent workshop hosted June 25-29, 2012, in St. Louis was attended by 60 people from laboratories around the world. Attendees have uniformly rated the workshop as outstanding on post-workshop web survey. We will continue to host biannual workshops to train new site managers and developers. In addition we will host, often in conjunction with scientific and technology conferences, “hackathon” events built around focused topics. An initial event on Pyxnat was hosted in St. Louis in Fall,

2011, and a second event on neuroimaging applications will be held in September, 2012, in association with the Neuroinformatics Congress in Munich, Germany. Future hackathon topics include cancer informatics, data visualization, and cloud computing services.

Data Mining. We have recently identified a new category of users that we have largely neglected to support up to this point. These users are technically savvy scientists who are not interested in developing new XNAT features per se but are interested in learning how to data mine XNAT repositories. We will develop specific documentation and training matreials for these users, with a focus on the REST API, data feeds, scripting tools, statistics engine, and computation systems. We believe that providing this class of users with rich tools and training will dramatically expand the impact of the scientific data hosted in XNAT repositories.

Outreach Coordinator. A full-time outreach coordinator with a strong programming and operations background will direct outreach and support activities, including production and maintenance of documentation and training materials and organization of the workshops, hackathons, and integration activities.

PROGRAM-WIDE ACTIVITIES

Software development practices. The XNAT team uses an agile development process that includes daily scrums, paired programming on high risk components, and continuous integration. Testing includes JUnit-based unit testing, automated functional testing, and automated penetration testing. We are currently developing a performance and scalability test suite that will be used to simulate heavy user loads. The test suite will be integrated into our continuous build system to verify that ongoing development does not degrade performance. All source code is managed through the Mercurial distributed version control system and is hosted for public access on Bitbucket. Development tasks and reported issues are managed using the Jira project tracking software. Several publicly accessible versions of XNAT are available to internal and external developers. Documentation is provided on the XNAT website and linked to NITRC.

Risk management. Despite the large scope and innovative nature of this project, the overall risk level is low, due to the experienced research team and our strategy deploying the software into real world projects early in the development process. Nonetheless, the project does carry some risks. First, the UI revision could lead to confusion for users if deployed midway into a project. To mitigate this, we will provide extensive training documentation specifically geared to users of the current UI. Second, the migration to Spring and Hibernate will require experienced XNAT developers to learn a new architecture and to migrate a fair amount of existing code. Again, we will work to minimize frustration by providing dedicated migration documentation and hosting a focused workshop to assist in the process. Finally, the clinical research profile is fairly high risk, as it is entering a new niche. Given the potential of translational imaging research, we believe this risk is well warranted. The Early Adopters Program was specifically created to identify critical issues early in the project and to adjust implementation strategies accordingly.

Timeline. XNAT releases will occur every 18 months, and packages of plugins will be released every.6 months (**Table 4**). Coordination with each of the evaluation projects described in Aim 2 will begin at the outset of the project and continue through the duration of the project. Outreach and support activities will be ongoing throughout the project.

Version	Release Date
Plugins #1	January, 2014
XNAT 1.7	July, 2014
Plugins #2	January, 2015
Plugins #3	July, 2015
XNAT 2.0	July, 2016
Plugins #4	January, 2017
Plugins #5	July, 2017
XNAT 2.1	January, 2018
Plugins #6	July, 2018

Table 4. XNAT Release schedule

CONCLUSION

This proposal continues our mission of creating exceptional software for the biomedical research community. The proposed deliverables are the result of discussion with the research community and will be developed, evaluated, and continuously improved through early deployment in real world projects. Our outreach plan ensures that the resulting software will be widely used to the benefit of a large community of scientists using imaging in their research.

References

1. Matthews, P.M., Rabiner, I., Gunn, R. (2011) Non-invasive imaging in experimental medicine for drug development. DOI: 10.1016/j.coph.2011.04.009. Curr Opin Pharmacol. 2011 Oct; 11(5): 501-7. PMCID: PMC Journal – in process
2. Davis, M., Hanson, S., Altevogt, B., *Rapporteurs*, Forum on Neuroscience and Nervous System Disorders, Institute of Medicine (2008) Neuroscience Biomarkers and Biosignatures: Converging Technologies, Emerging Partnerships: Workshop Summary. National Academies Press.
3. Zerhouni, E., 2008. "NIBIB's Fifth Anniversary Symposium : changing the world's healthcare through biomedical technologies." Keynote lecture. Available at http://videocast.nih.gov/Summary.asp?File=13863.
4. Das, S., Zijdenbos, A. P., Harlap, J., Vins, D., & Evans, A. C. (2011). LORIS: a web-based data management system for multi-center studies. DOI:10.3389/fninf.2011.00037. Front Neuroinform. 2011; 5: 37. PMCID: PMC3262165
5. Freymann, J. B., Kirby, J. S., Perry, J. H., Clunie, D. A., & Jaffe, C. C. (2012). Image data sharing for biomedical research--meeting HIPAA requirements for De-identification. DOI:10.1007/s10278-011-9422-x. J Digit Imaging. 2012 Feb; 25(1): 14–24. PMCID: PMC3264712
6. Ozyurt, I. B., Keator, D. B., Wei, D., Fennema-Notestine, C., Pease, K. R., Bockholt, J., & Grethe, J. S. (2010). Federated web-accessible clinical data management within an extensible neuroimaging database. DOI:10.1007/s12021-010-9078-6. Neuroinformatics. 2010 Dec; 8(4) 231-49. PMCID: PMC2974931
7. Scott, A., Courtney, W., Wood, D., de la Garza, R., Lane, S., King, M., Wang, R., et al. (2011). COINS: An Innovative Informatics and Neuroimaging Tool Suite Built for Large Heterogeneous Datasets. DOI: 10.3389/fninf.2011.00033. Front Neuroinform. 2011; 5: 33. PMCID: PMC3250631
8. Van Horn, J. D., & Toga, A. W. (2009). Is it time to re-prioritize neuroimaging databases and digital repositories? DOI:10.1016/j.neuroimage.2009.03.086. Neuroimage. 2009 Oct 1; 47(4): 1720-34. PMCID: PMC2754579
9. Archie, K. A., & Marcus, D. S. (2012). DicomBrowser: Software for Viewing and Modifying DICOM Metadata. DOI:10.1007/s10278-012-9462-x. J Digit Imaging. 2012 Feb 15. [Epub ahead of print]
10. OWASP Top Ten Vulnerabilities Reference available at https://www.owasp.org/index.php/Category:OWASP_Top_Ten_Project.
11. Fielding, R.T.,2000. "Architectural Styles and the Design of Network-based Software Architectures" PhD Thesis, University of California, Irvine. Accessible at http://www.ics.uci.edu/~fielding/pubs/dissertation/top.htm.
12. REST API Reference Documentation, available at https://wiki.xnat.org/pages/viewpage.action?pageId=5013754.
13. W3C Provenance Reference Documentation available at http://www.w3.org/2011/prov/wiki/Main_Page.
14. XNAT 2.0 User Interface Specification available at https://wiki.xnat.org/download/attachments/5016988/FRD+Final+-+Deliver.docx?version=1&modificationDate=1340806591373.
15. XNAT 2.0 Reference Page Designs available at http://integrity.prevuu.com/p/xnat.
16. Doran S.J., d'Arcy J., Collins D.J., Andriantsimiavona R., Orton M., Koh D., Leach, M.O. (2012) Development of a research PACS for functional imaging analysis in clinical research and clinical trials. Radiographics, in press.
17. Ashish, N., Ambite, J. L., Muslea, M., & Turner, J. A. (2010). Neuroscience Data Integration through Mediation: An (F)BIRN Case Study. DOI:10.3389/fninf.2010.00118. Front Neuroinform. 2010 Dec 28; 4:118. PMCID: PMC3017358
18. Gerhard, S., Daducci, A., Lemkaddem, A., Meuli, R., Thiran, J.-P., & Hagmann, P. (2011). The connectome viewer toolkit: an open source framework to manage, analyze, and visualize connectomes. DOI:10.3389/fninf.2011.00003. Front Neuroinform. 2011; 5:3. PMCID: PMC3112315
19. Marcus, D. S., Harwell, J., Olsen, T., Hodge, M., Glasser, M. F., Prior, F., Jenkinson, M., et al. (2011). Informatics and data mining tools and strategies for the human connectome project. DOI:10.3389/fninf.2011.00004. Front Neuroinform. 2011; 5:4. PMCID: PMC3127103
20. Schwartz, Y., Barbot, A., Thyreau, B., Frouin, V., Varoquaux, G., Siram, A., Marcus, D. S., et al. (2012). PyXNAT: XNAT in Python. DOI:10.3389/fninf.2012.00012. Front Neuroinform. 2012; 6:12. PMCID: PMC3354345
21. Reigel, T., Hennessey, J., Schneider, B., Brown, T., Kolasny, A., Trinh, H. Wang, L. (2011) Region of interest shape analysis using computational anatomy works. Program No. 841.05 Neuroscience Meeting Planner. Washington, DC: Society for Neuroscience, 2011. Online.

22. Poline, J.-B., Breeze, J. L., Ghosh, S., Gorgolewski, K., Halchenko, Y. O., Hanke, M., Haselgrove, C., et al. (2012). Data sharing in neuroimaging research. DOI:10.3389/fninf.2012.00009. Front Neuroinform. 2012; 6:9. PMCID: PMC Journal – in process
23. Freymann, J. B., Kirby, J. S., Perry, J. H., Clunie, D. A., & Jaffe, C. C. (2012). Image data sharing for biomedical research--meeting HIPAA requirements for De-identification. DOI:10.1007/s10278-011-9422-x. J Digit Imaging. 2012 Feb; 25(1): 14-24. PMCID: PMC3264712 [Available on 2013/2/1].
24. Van Essen, D. C., Ugurbil, K., Auerbach, E., Barch, D., Behrens, T. E. J., Bucholz, R., Chang, A., et al. (2012). The Human Connectome Project: A data acquisition perspective. DOI:10.1016/j.neuroimage.2012.02.018. Neuroimage. 2012 Feb 17. [Epub ahead of print]
25. Marcus, D. S., Harwell, J., Olsen, T., Hodge, M., Glasser, M. F., Prior, F., Jenkinson, M., et al. (2011). Informatics and data mining tools and strategies for the human connectome project. DOI:10.3389/fninf.2011.00004. Front Neuroinform. 2011; 5:4. PMCID: PMC3127103

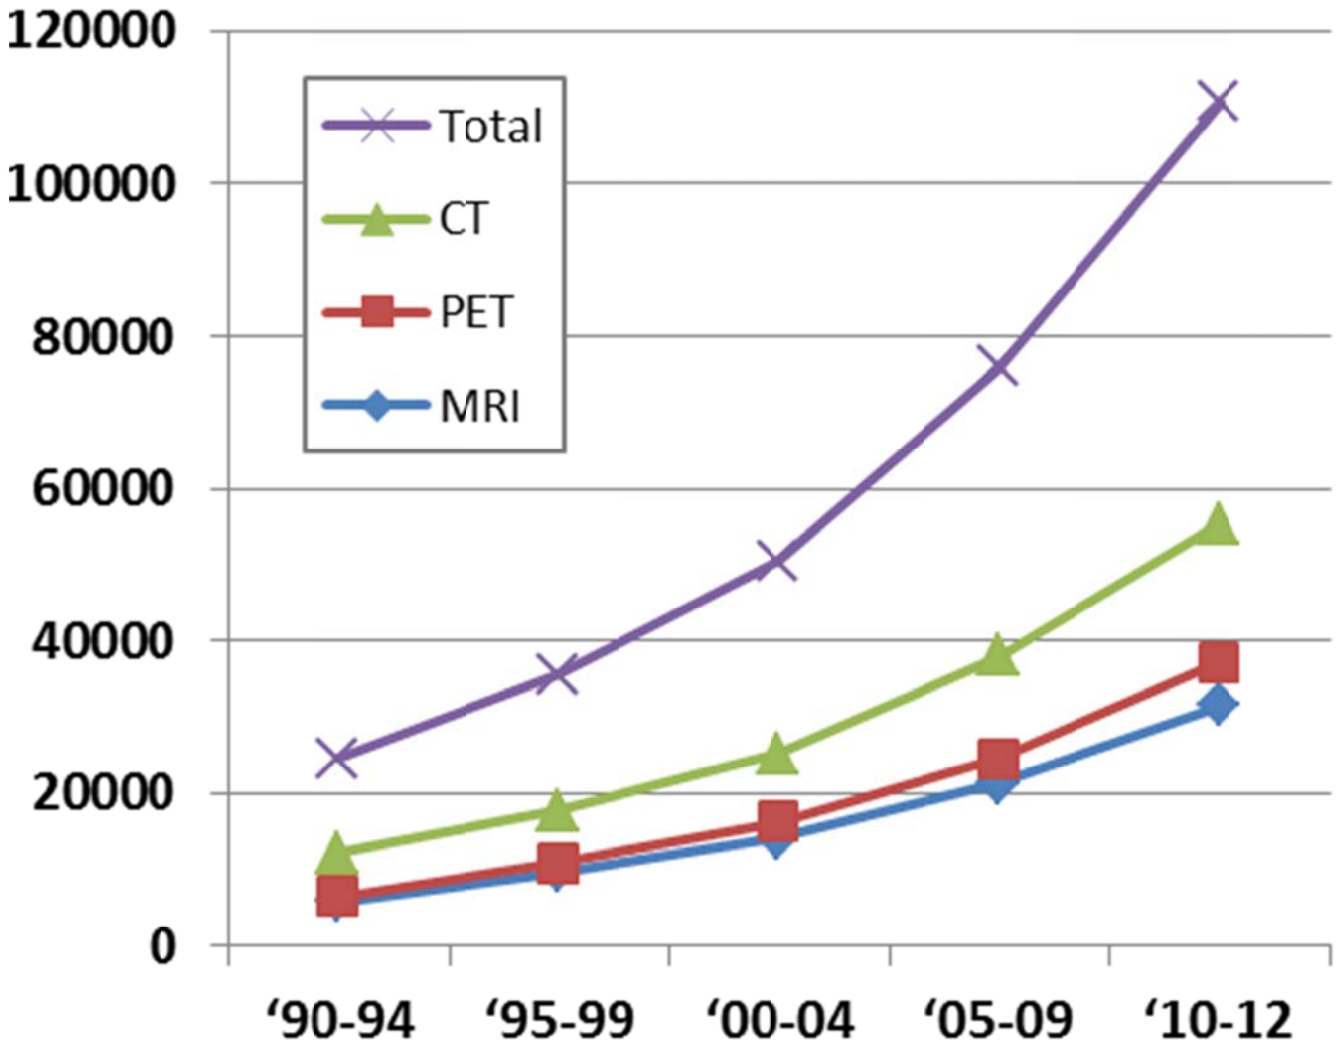
120000
100000
80000
60000
40000
20000
0
Total
CT
PET
MRI
'90-94
'95-99
'00-04
'05-09
'10-12

Figure 1

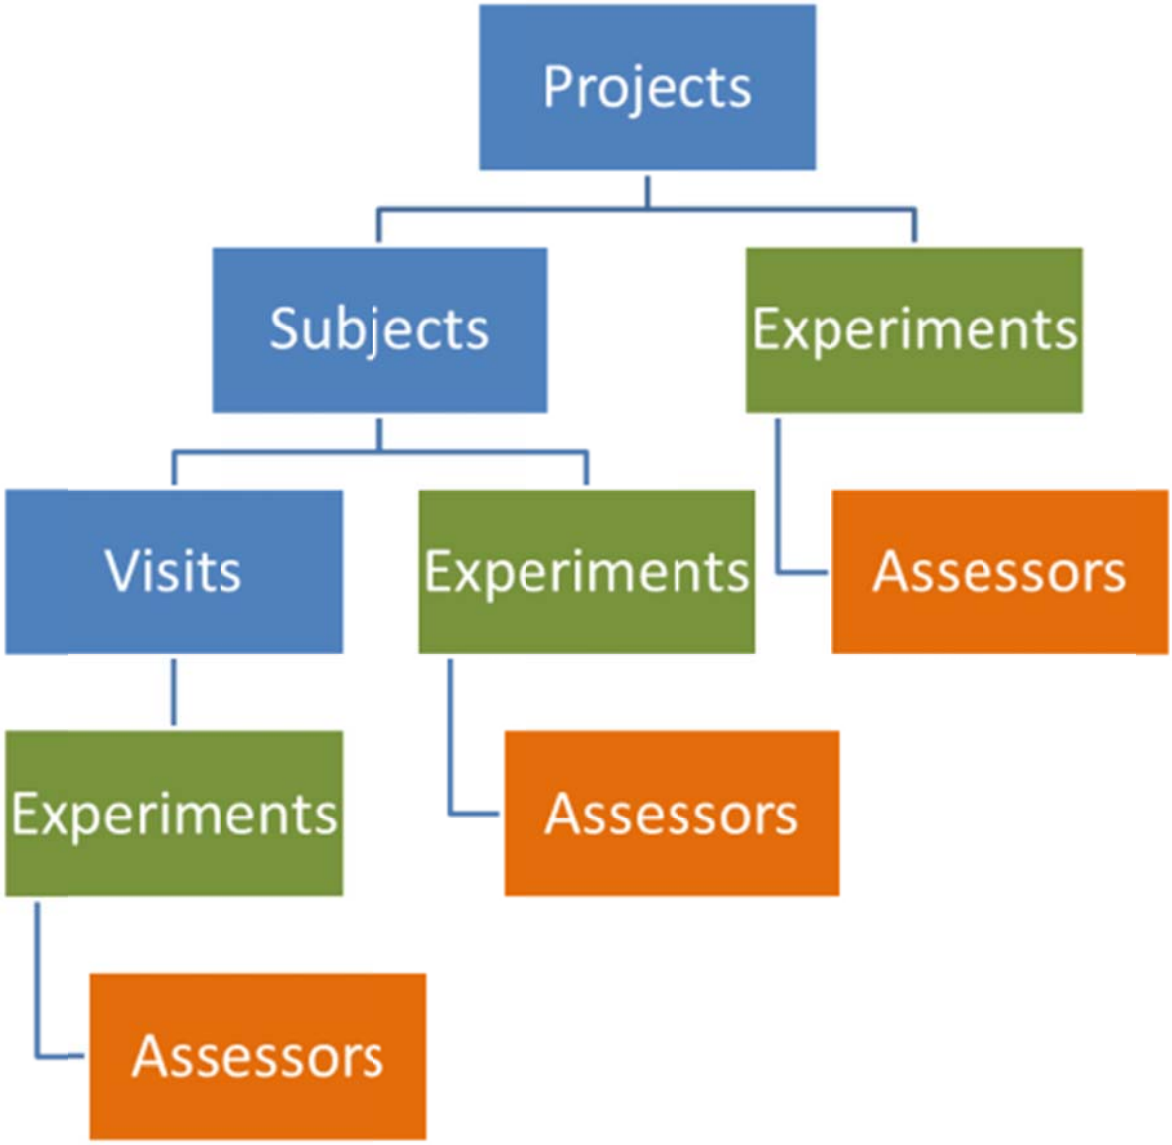
Projects
Subjects
Experiments
Visits
Experiments
Assessors
Experiments
Assessors
Assessors

Figure 2

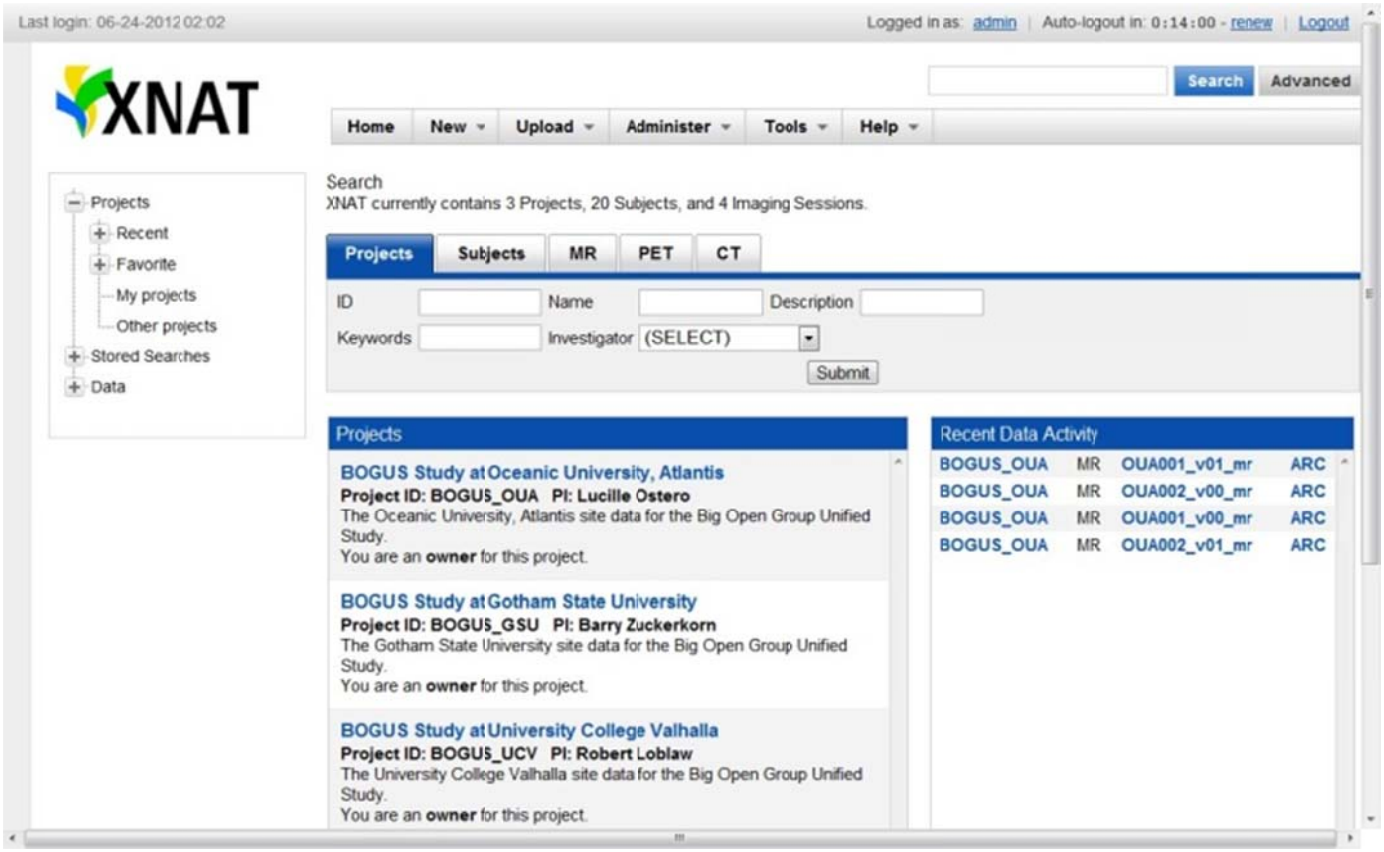

Figure 3, left

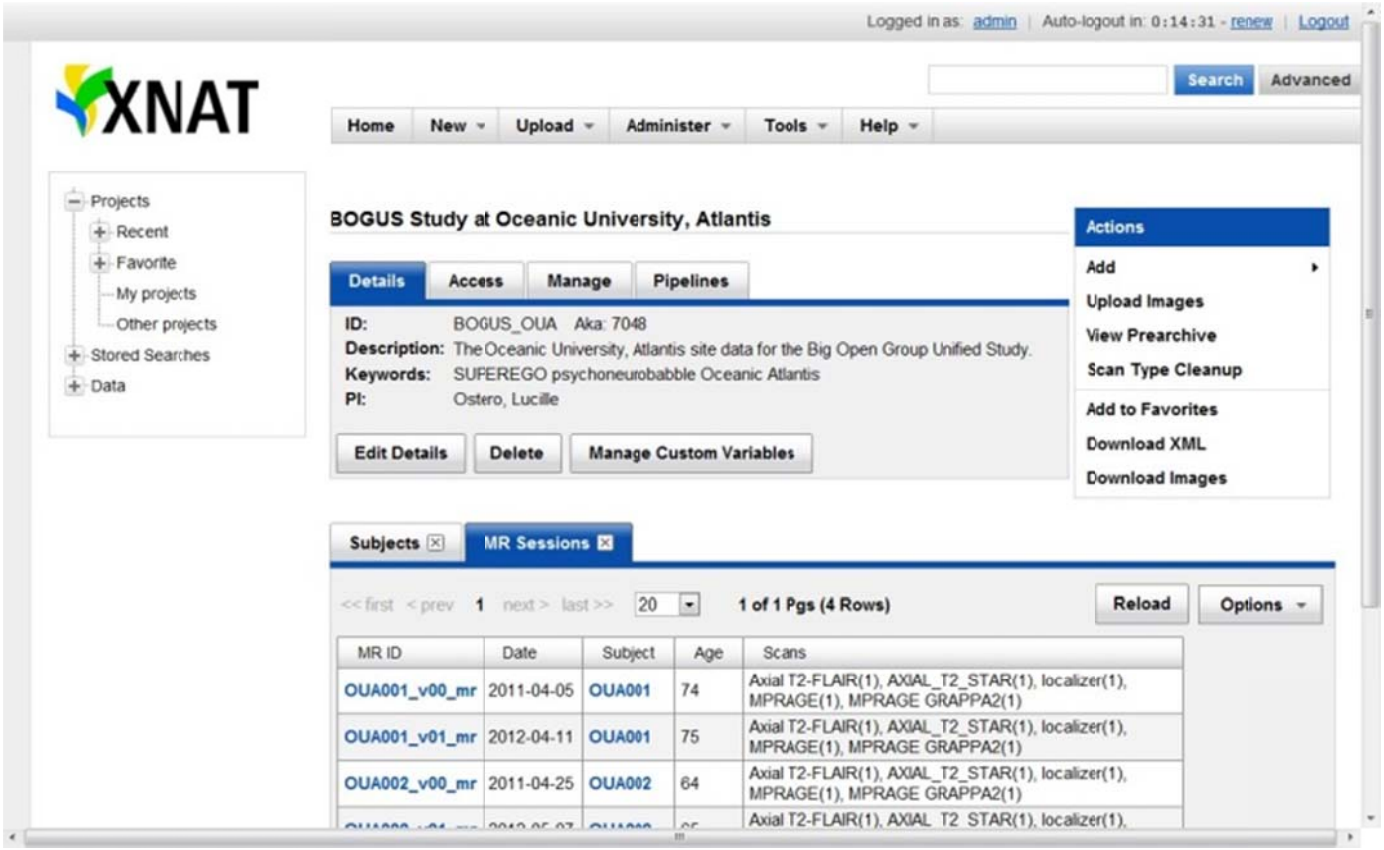

Figure 3, center

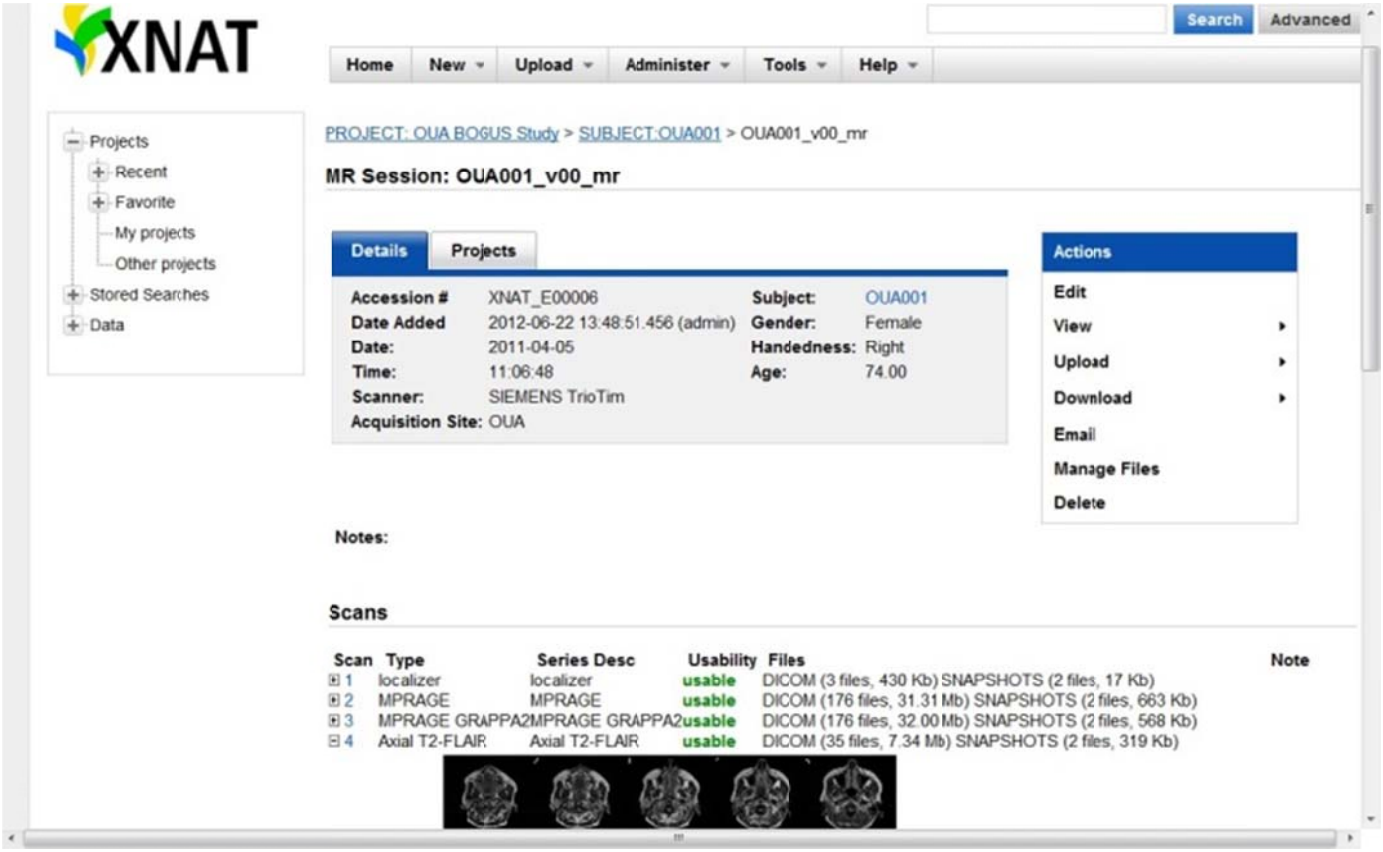

Figure 3, right

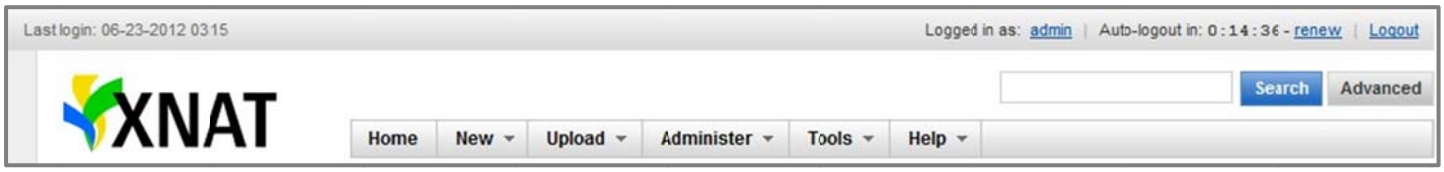

Figure 4

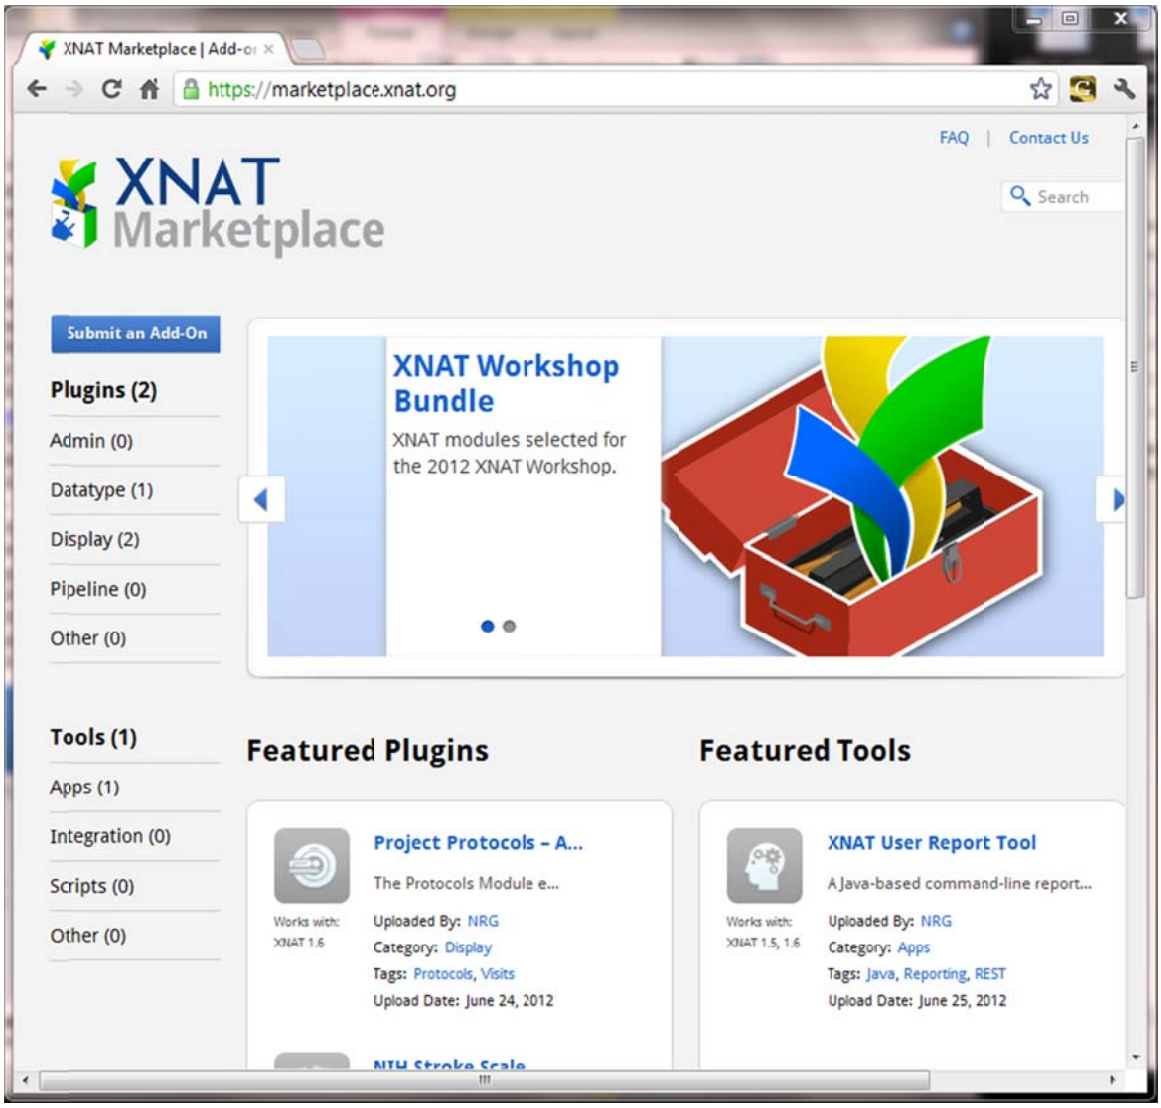

Figure 5, left

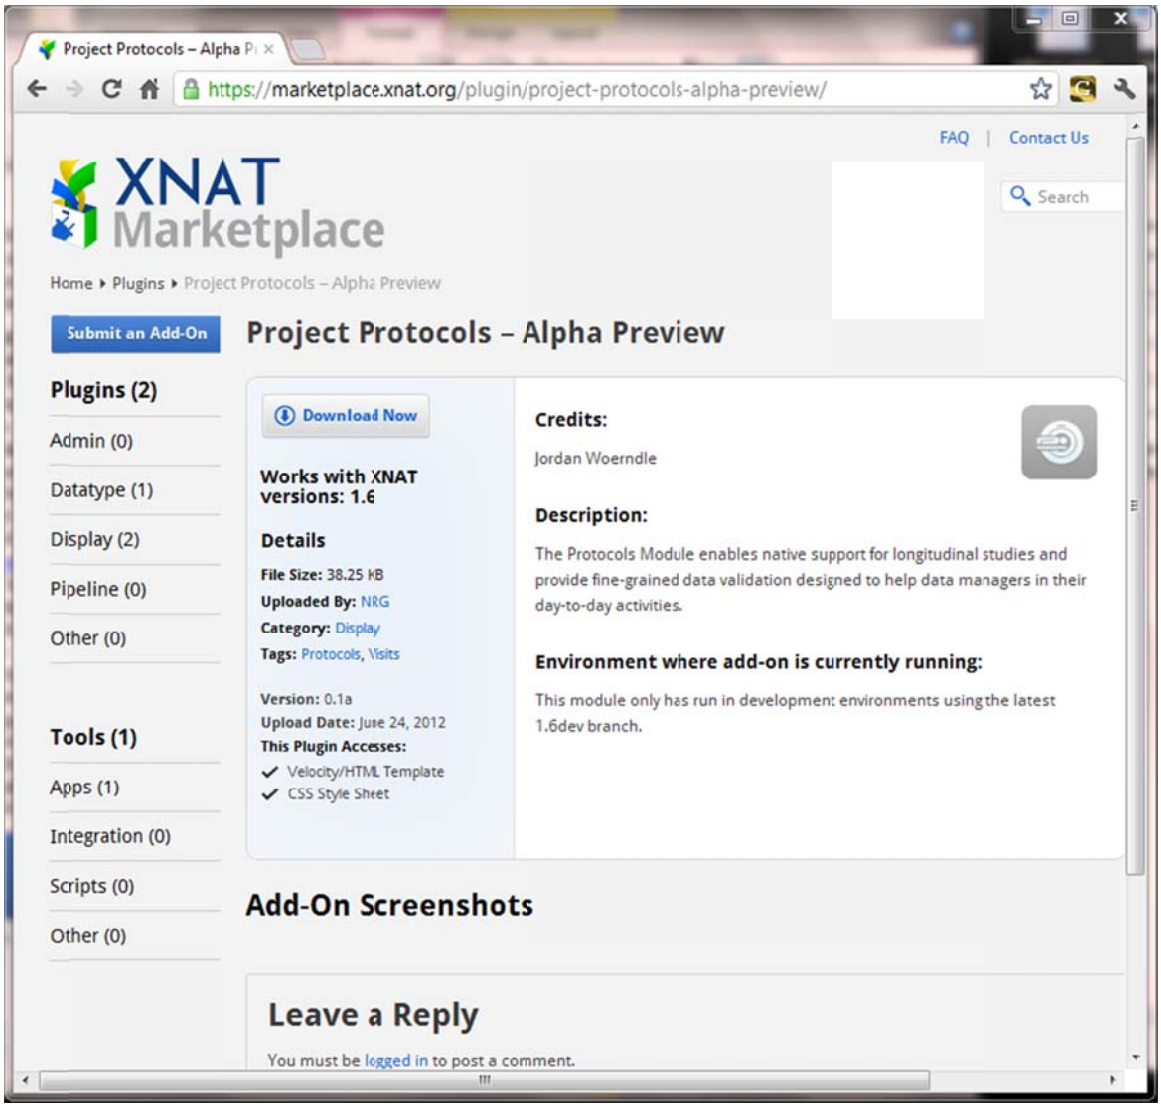

Figure 5, right

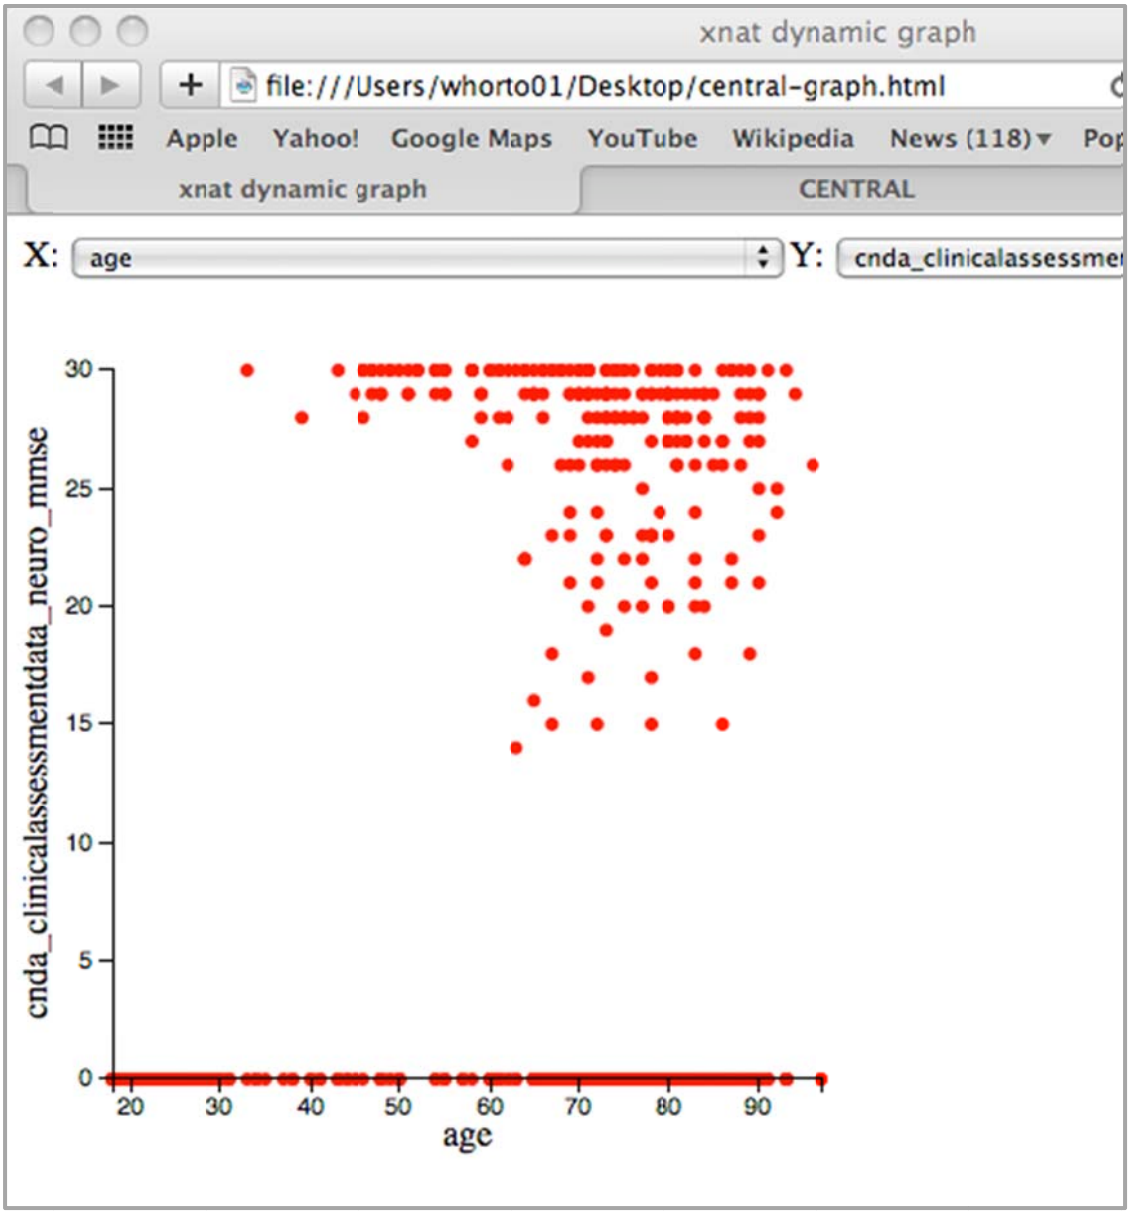
xnat dynamic graph
file:///Users/whorto01/Desktop/central-graph.html
Apple
Yahoo!
Google Maps
YouTube
Wikipedia
News (118)
xnat dynamic graph
CENTRAL
X:
age
Y:
cnda_clinicalassessmentdata_neuro_mmse
30
25
20
15
10
5
0
20
30
40
50
60
70
80
90
age

Figure 6

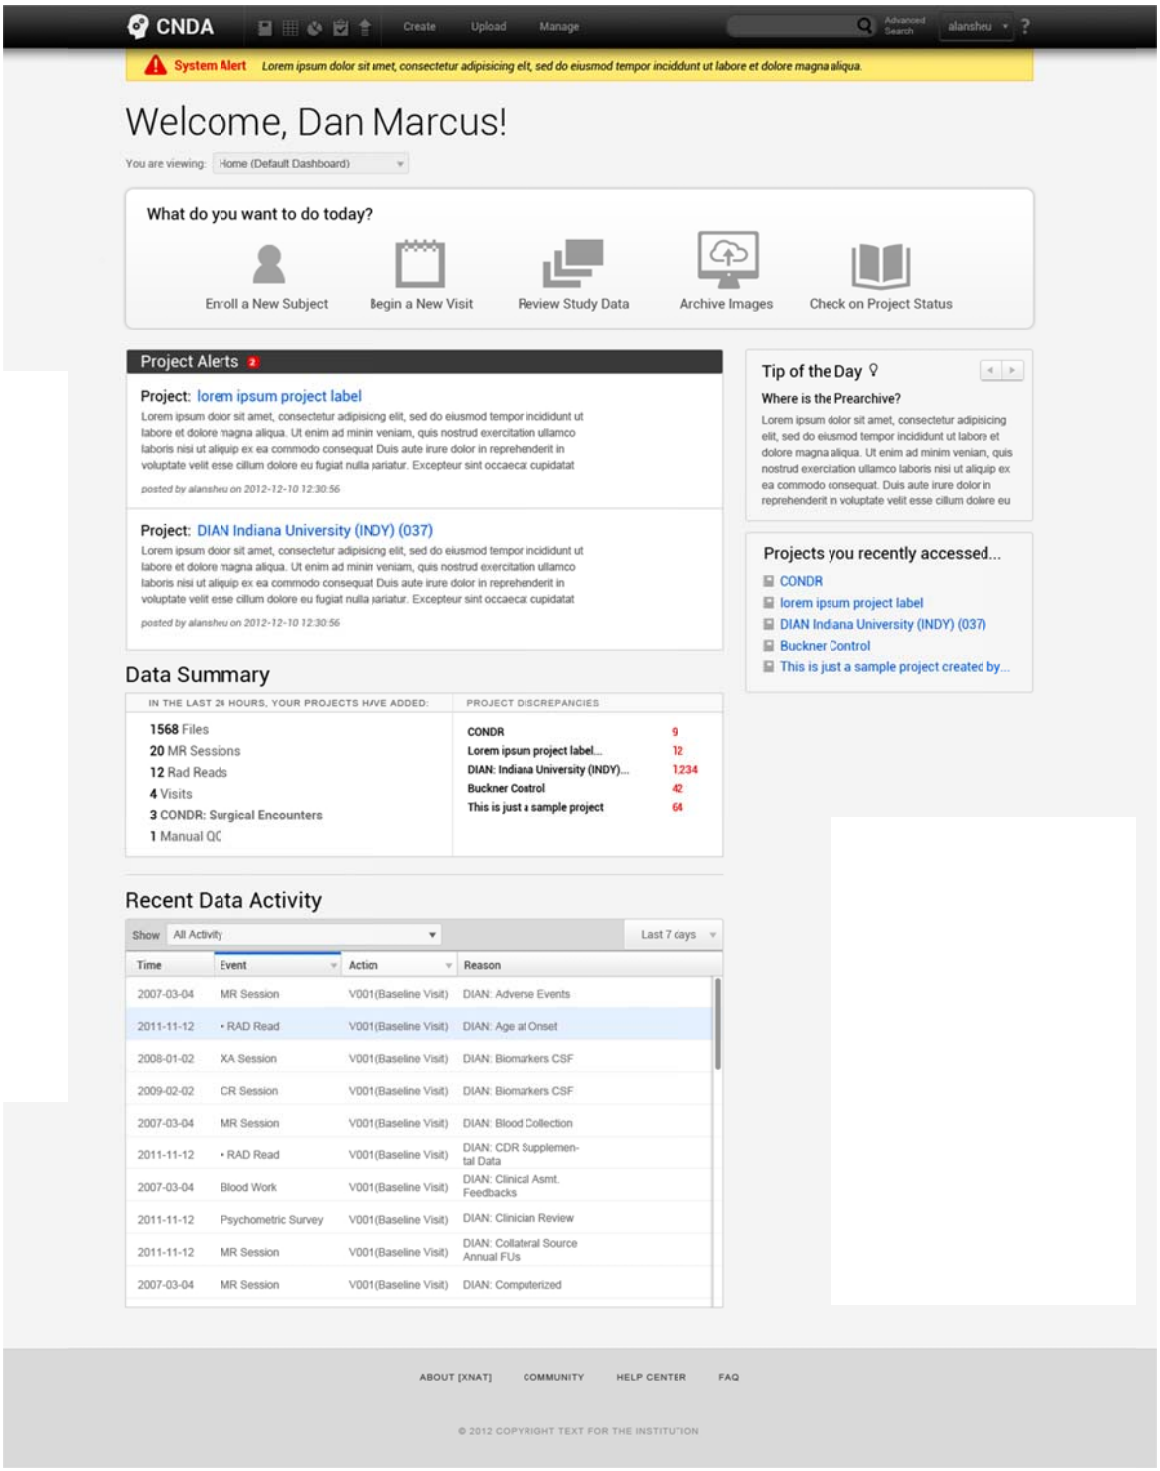

Figure 7, far left

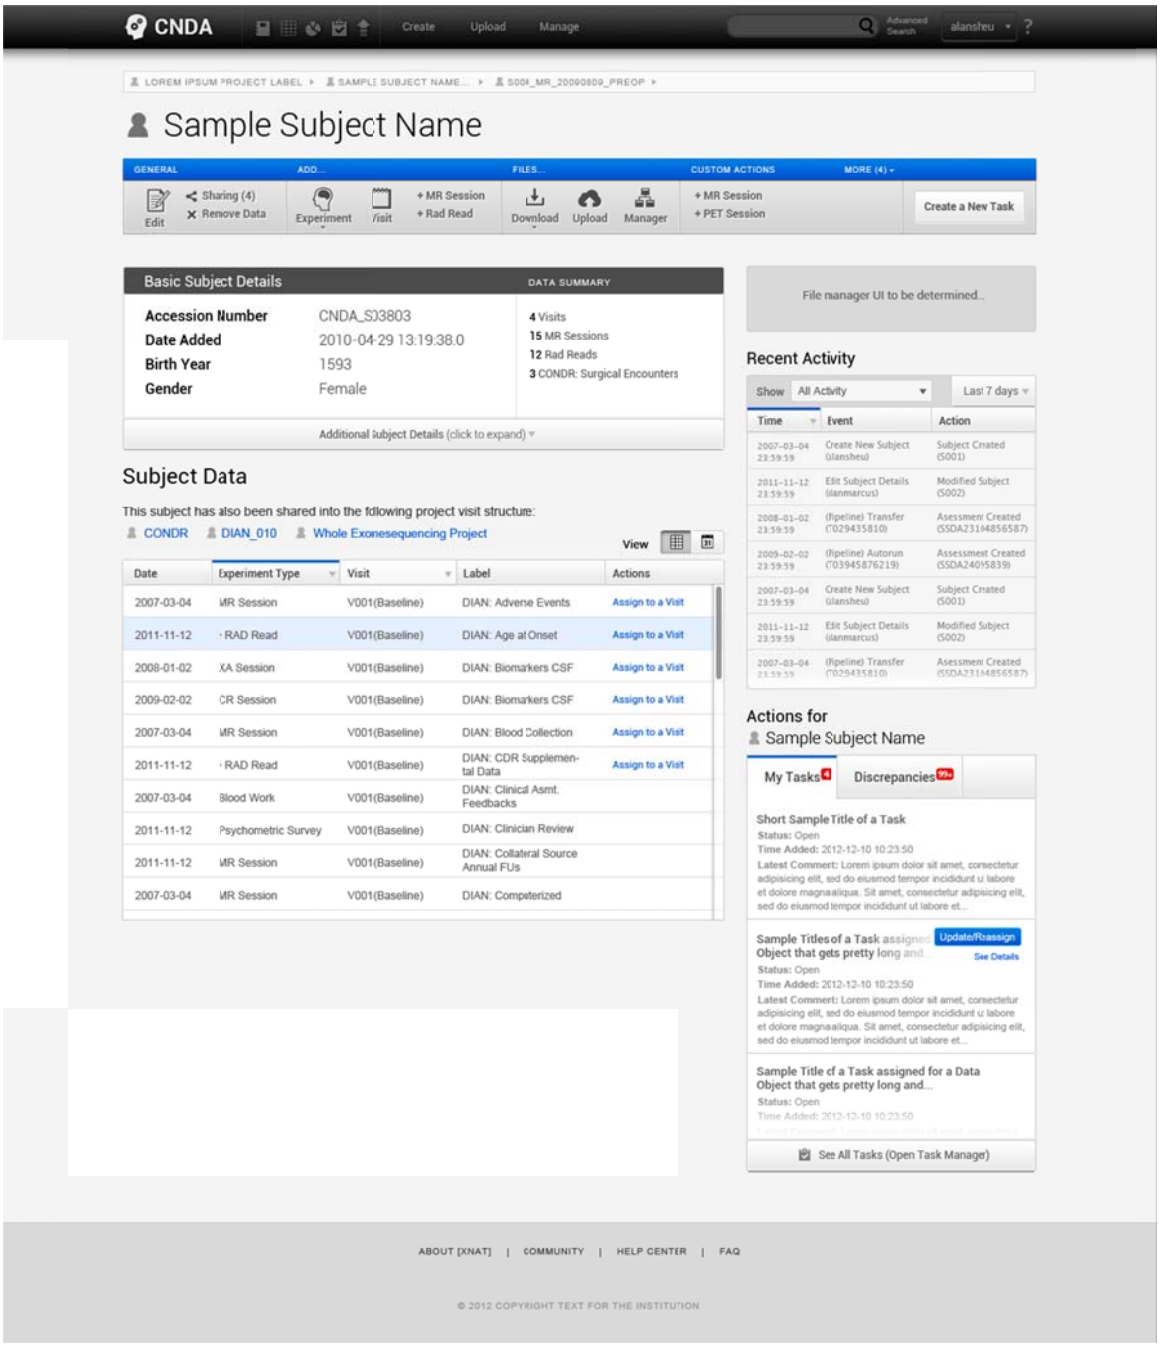

Figure 7, left

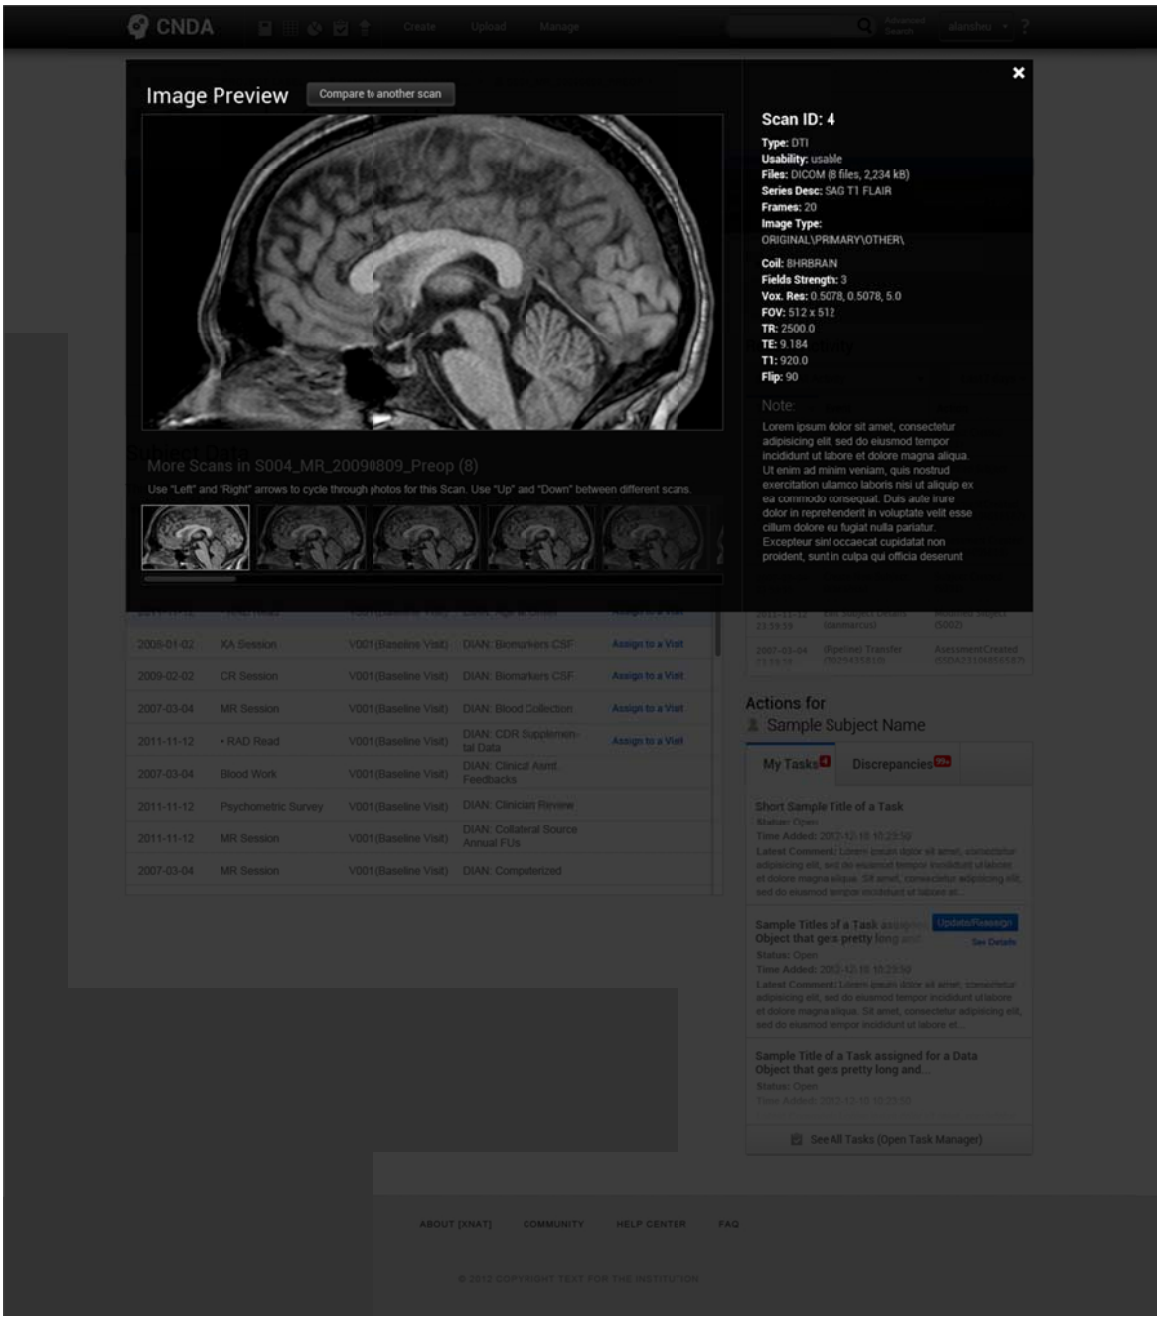

Figure 7, right

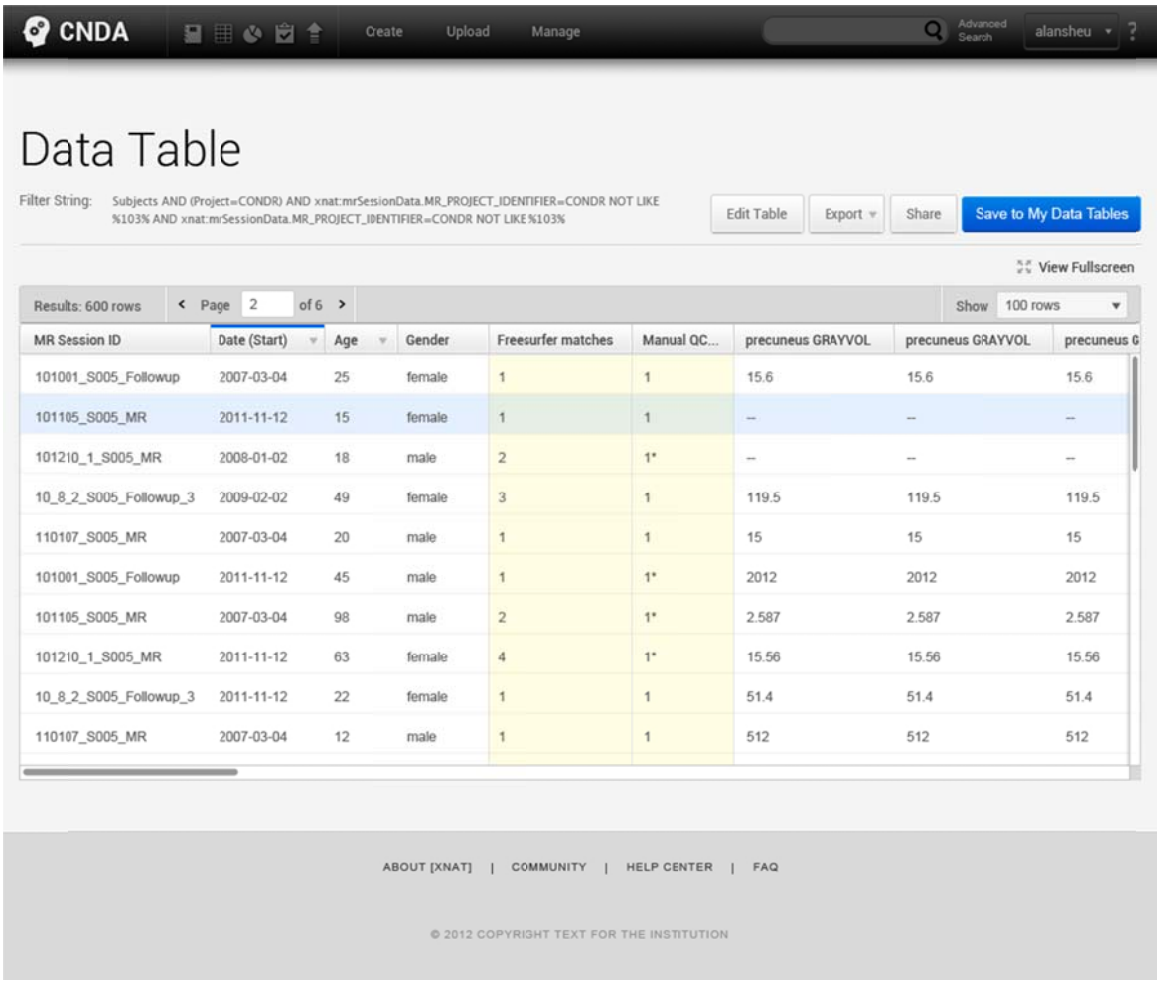

Figure 7, far right