

# Software Tools for Analysis and Visualization of FMRI Data

*NMR in Biomedicine*, in press.

Robert W. Cox and James S. Hyde

Biophysics Research Institute  
Medical College of Wisconsin  
8701 Watertown Plank Road  
Milwaukee, WI 53226-0509

e-mail: [rwcox@mcw.edu](mailto:rwcox@mcw.edu)

Contact: Robert W. Cox

Address: (as above)

Phone: 414-456-4038

FAX: 414-266-8515

**Running Title:** Software Tools for FMRI

**Abbreviations used:**

ANOVA, Analysis of Variance;  
ASCII, American Standard Code for Information Interchange;  
BOLD, Blood Oxygenation Level Dependent;  
CD-ROM, Compact Disk Read Only Memory;  
CNS, Central Nervous System;  
DICOM, Digital Imaging and Communication in Medicine;  
FMRI, Functional Magnetic Resonance Imaging;  
GNU, Gnu's Not Unix;  
HDF, Heirarchical Data Format;  
HTML, Hypertext Markup Language;  
HTTP, Hypertext Transport Protocol;  
I/O, Input/Output;  
MCW, Medical College of Wisconsin;  
NCSA, National Center for Supercomputing Applications;  
PET, Positron Emission Tomography;  
RMS, Root Mean Square;  
ROI, Region of Interest;  
SPGR, Spoiled Gradient Recalled at Steady State;  
VRML, Virtual Reality Markup Language;  
WWW, World Wide Web.

---

## Introduction

---

The amount of data gathered with echo-planar BOLD-weighted functional MRI (fMRI) can be staggering. With commercially available equipment, it is routine to gather 10  $64 \times 64$  images per second almost continually for an hour (scanner: GE Signa 1.5 T, GE Medical Systems, Waukesha; local head coil: Medical Advances, Milwaukee). After reconstruction to 16 bit images, the data to be analyzed are on the order of 250 megabytes per hour. Faster scanning and higher resolution are available with custom-built equipment, yielding data rates in excess of 1 gigabyte per scanning hour.

Only a small portion of such a vast accumulation of numbers relates to neural activity. Most of the content is due to the baseline MR signal; after that is removed, much of the variance is related to the cardiac and respiratory cycles, or to subject<sup>1</sup> and scanner instabilities. As far as neurological investigations are concerned, these portions of the signal are ‘noise’. (But see Biswal *et al.*,<sup>2</sup> for a deeper investigation of this issue.)

Static MR imaging methods provide many techniques for probing the structure and function of human subjects and patients. The addition of the time dimension only increases the flexibility of this instrument. As new methods are invented for dynamic imaging of the CNS, new analytical methods will be needed; nor can existing methods for analysis of fMRI time series be considered closed to further improvement.

Analysis of complex multidimensional data sets includes display of the raw data and derived quantities; visualization can be considered part of analysis, where the investigator’s pattern recognition skills are coupled to the computer’s imaging capabilities.

It is still relatively difficult to start an FMRI program at an institution with no experience. There are three types of expertise required: (a) MR physics or engineering, since FMRI pulse sequences push current scanner hardware rather severely; (b) Statistics and software development, since the large quantities of data must be analyzed in many different ways; and (c) Neuropsychology, since the results of the analyses must be interpreted in the light of existing knowledge about the brain. Attempting to start FMRI research without the first two classes of skills present will be a frustrating experience.

In this paper, we discuss the software tools needed for FMRI, and in particular one toolset that is freely available: *AFNI* from the Medical College of Wisconsin (MCW)<sup>3</sup>. This package has been under development since mid-1994, and now comprises over 75,000 lines of C, running under Unix<sup>TM</sup> and X11 Windows. The heart of the package is the `afni` program itself (about 20,000 lines); in addition, there are over 25 auxiliary programs for manipulation of FMRI data sets. One of our goals with the release of *AFNI* outside MCW is to provide some of the processing expertise (b) to new FMRI sites. Another goal is to provide a basis for the sharing of interactive FMRI analysis tools, through the ‘plugin’ capability, which allows external C functions to be incorporated into `afni` at run time and to be executed by the user from the program’s graphical interface.

Many sites doing FMRI (and other functional neuroimaging methods, such as PET) are developing software systems for data visualization, analysis, and reduction. Many of these efforts are complementary in intention, but incompatible in implementation. Perhaps it is too late, but another purpose of this paper is to call for some coordination of FMRI tool development at research institutions.

We outline the structure, merits, and major gaps in *AFNI* in part to illustrate the magnitude and potential directions of the development task needed to provide a comprehensive fMRI visualization, analysis, and integration package to the research community. The remaining sections are organized around the central functions which such a tool should provide: data storage, interactivity, visualization, spatial normalization, analysis, integration, and package extensibility. At the end, we discuss what is needed for the fMRI community to create jointly such a comprehensive tool.

---

## Data Storage

---

Storage of image data is very far from being standardized. Over 100 file formats are listed as being ‘standards’ on the Internet—and this is just for 2D images. Every programmer and every manufacturer seems compelled to invent a new format for their particular and peculiar data.

*AFNI*. The fundamental unit of data storage is a 3D array, whose elements can be 8 bit integers, 16 bit integers, 32 bit floating point numbers, or 64 bit complex numbers. An *AFNI* ‘dataset’ comprises a file containing one or more such 3D arrays (the ‘brick’ file, with only image data), plus another file containing auxiliary information about the data (the ‘header’ file). Each element of the header consists of a unique identifying name and an associated array of integers, floats, or characters, all stored in ASCII format. (One example of header information are the entries giving the spatial size and location of the 3D arrays in the brick file. Time-dependent dataset headers contain information about the temporal spacing and ordering of the slices within the 3D arrays.) When a dataset is input, the *AFNI* programs search the auxiliary information using the identifying name; auxiliary arrays with names that are not needed by the particular program are simply ignored. In this way, new

types of information can be added to the header file without causing any existing programs to fail.

Storing large data sets as big arrays with associated named data attributes is not unique to *AFNI*. The ideas above were adapted from the HDF (hierarchical data file) format from NCSA. Originally, *AFNI* used the HDF format, but the input/output (I/O) overhead was too cumbersome for an interactive program: more than a factor of 10 slower than directly programmed I/O for a 16 megabyte file, for example.

The *AFNI* data format is designed to be good at storage of rectangularly sampled volumetric data. It is relatively simple to use such datasets with other programs, since the image data is stored separately in one big file. If necessary, another program can ignore the *AFNI* header file and read data directly from the brick file, which has essentially no structure to comprehend.

*Desiderata.* The *AFNI* data format could be used to store rectangular scan  $k$ -space FMRI raw data, although we do not use it for that purpose at MCW. It could also be adapted to store non-rectangular scan data by storing a description of the  $k$ -space path in the first data brick, followed by the data samples from the actual imaging sequences in later data bricks.

The relative ease of gathering the data, and the variability in detected activation between scanning runs (even in the same subject), mean that multiple runs and multiple subjects are usually gathered in any neuroscience investigation with FMRI. It is therefore desirable to link individual FMRI data sets together into larger associations, and to maintain computerized records of the relationships in these collections. When derived data sets are created (*e.g.*, by averaging), they should be entered into the database of data sets, along with the description of the methods used to create them.

The *AFNI* data format contains no facility for describing complex geometric data, such as surfaces (*e.g.*, the gray-white matter interface) or irregular anatomical regions-of-interest (ROIs). In the future, we plan to incorporate a subset of the VRML (Virtual Reality Markup Language) standard to describe complex objects embedded within 3D bricks. This standard, developed for WWW visualization purposes, contains facilities for describing points, curves, surfaces, and solids. Although VRML is not ideal for application to 3D and 4D medical images, there are many points in its favor: (a) it is an international standard; (b) it is machine and resolution independent; (c) some freely available software libraries exist for interpreting VRML files and for rendering VRML defined objects into displayable images; and (d) VRML-capable WWW browsers are available and can act as quick viewers for VMRL code generated by FMRI software.

---

## **Interaction**

---

The modern computer user likes everything to be interactive. Besides being gratifying, easily used interaction can help the user explore his data. Given the complex structure of FMRI data sets, exploration and quick trial analyses should be encouraged.

Not all functions are easily made interactive. With present technology, integrating the results from 15 subjects in 8 scan conditions each, involving over 2 gigabytes of raw image data, takes several minutes at best. Exploring the parameters of such a lengthy integrative operation is not interactively feasible. Instead, a number of runs needs to be made in batch mode, perhaps overnight, and the results visualized interactively later.

AFNI. Only some functions are interactive in the *AFNI* package. Others have been developed in batch-only programs, which are usually run from scripts. Part of the reason for this state of affairs is that integration of new operations into the interactive program `afni` used to be relatively difficult and tedious. With the recent introduction of the plugin extension feature to `afni`, more capabilities will be made available in both interactive and batch modes.

Besides visualization, the following functions are interactive in the `afni` program: computation of functional activation using the correlation method,<sup>8, 4</sup> adjustment of statistical thresholds for activation detection, resampling to a different grid spacing, and transformation to Talairach (stereotaxic) coordinates.<sup>9</sup> An interactive plugin has been developed to perform various editing tasks on 3D datasets (*e.g.*, clustering of active voxels). Three plugins that operate on voxel time series do linear least squares fits, simple statistics, and Fourier transforms.

Functions that are not yet interactive include image registration and all operations that involve more than one dataset. An example of the latter are the 3dANOVA programs for performing voxel-by-voxel analysis of variance on many different 3D datasets.

Desiderata. Most operations should be available in batch and interactive forms. Batch mode is very useful when performing routinized forms of analysis and data reduction (*e.g.*, image reconstruction is usually performed in batch mode at MCW). A set of utilities that perform basic FMRI tasks is very useful, and can be assembled into scripts in many ways. It is important to remember that putting together such scripts is really a form of programming, and so will be difficult for many users. For this purpose, a set of sample scripts, or a script generation program, would be very useful. Nothing like this has been developed at MCW as yet.



One path that leads easily to paired batch and interactive capabilities is the casting of each operation into the form a set of library functions. A batch program then consists of an interface program that reads the command line and appropriately calls the library; an interactive interface needs to create a window for the user to enter the appropriate parameters for the function calls. By describing the library operations appropriately, the batch and interactive interfaces could in principle be generated directly, instead of being coded manually. At MCW, we are beginning to move in this direction with the *AFNI* plugin capability, in which the programmer provides a specification of the interface, and `afni` then generates a graphical interface window that matches. We plan to extend this capability to enable plugins to have a batch interface shell generated from the same interface specification, thus allowing simultaneous development of batch and interactive tools.

---

## Visualization

---

Techniques for viewing complex multidimensional data sets are an active area of research, involving computer science, software and hardware engineering, and visual psychophysics.

*AFNI*. Program `afni` is capable of displaying orthogonal slices from 3D bricks, as shown in Fig. 1. The number of slices in each image window is user controlled. The orthogonal views are linked together at a given viewpoint (where the crosshairs intersect), whose stereotaxic coordinates are displayed in the control panel. Clicking the mouse on any image will cause the viewpoint of all windows to jump to that location.

Time-dependent datasets can also be graphed, with the central voxel in the graph array linked to the image viewpoint. In Fig. 1, the central time series graph also shows a smooth waveform: this is the correlation reference.<sup>8</sup>

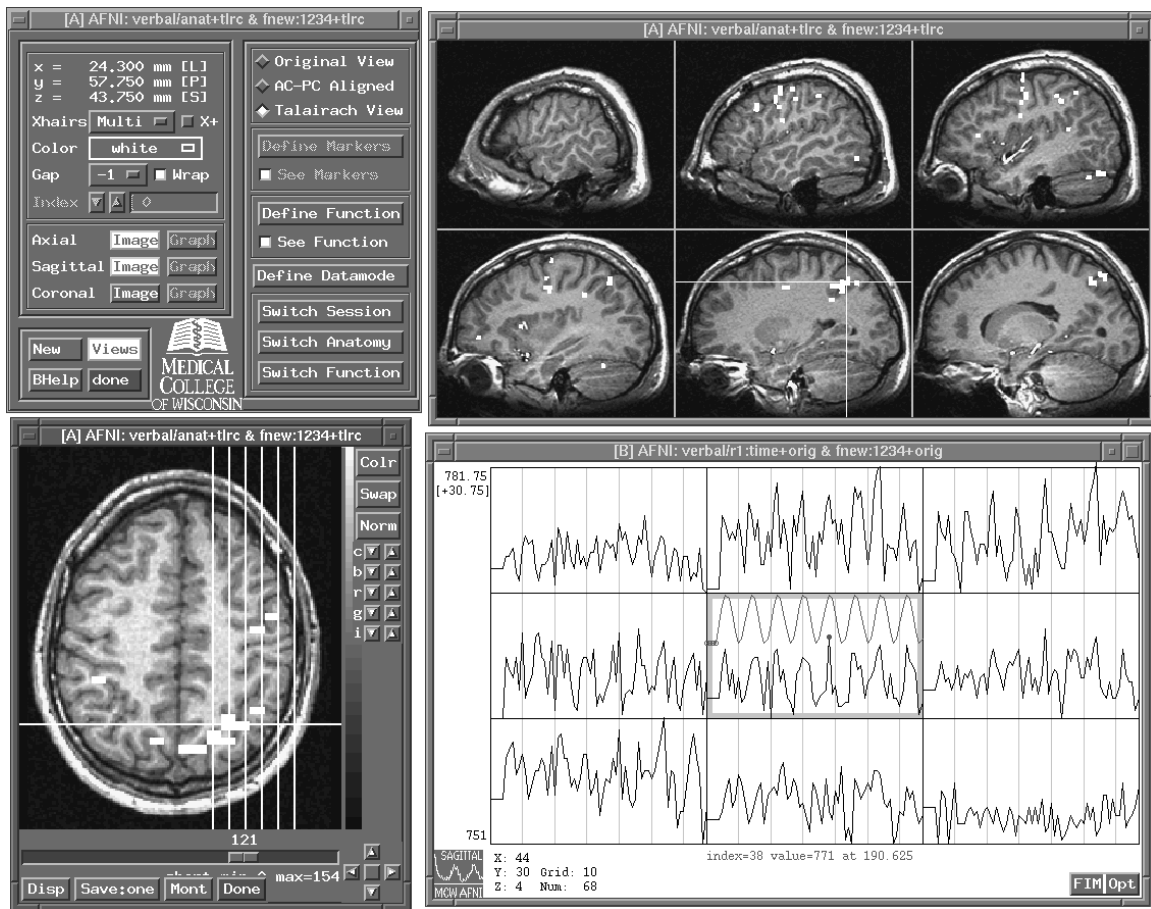


Figure 1. Screen shot from *AFNI*; features are discussed in the text.

One feature that distinguishes *afni* from general 3D brick visualization tools is the ability to overlay functional results gathered at one resolution onto anatomical data gathered at a higher resolution. Figure 1 shows functional activation at  $3.75 \times 3.75 \times 7.00 \text{ mm}^3$  resolution overlaid on 3D SPGR data gathered at  $0.94 \times 0.94 \times 1.20 \text{ mm}^3$ , with both resampled to stereotaxic coordinates on a  $1 \times 1 \times 1 \text{ mm}^3$  grid. In any image window, the user can view one anatomical dataset (as the grayscale underlay) and one functional dataset (as the color overlay). *Afni* can load multiple datasets from disk, and allows the user to switch among them as desired. Multiple sets of viewing windows can be opened to display datasets side-by-side; the viewpoints

of these different datasets can be chosen independently or locked together. The combination of these features makes it possible to scroll purposefully through vast amounts of data with relative ease.

*Desiderata.* Volume rendering and animation are two commonly used visualization tools that are missing from *AFNI*. At present, we have no plans to incorporate these features into our software, since so many other packages for these functions already exist.

At present, *AFNI* can only overlay two datasets if they are sampled on parallel grids; that is, if the acquisition planes are parallel or orthogonal. For some FMRI acquisitions, the functional slices are chosen to align with given anatomical features after the high resolution images are acquired. For this reason, it is desirable to have the ability to overlay datasets sampled on grids whose relative orientation is arbitrary. This feature will be necessary in any comprehensive FMRI tool.

Navigating through brain images can be quite overwhelming. A simplified visual guide would be very useful; in other words, a brain atlas linked to the data visualization tools. Commercial brain atlases are available on CD-ROM (*e.g.* VOXEL-MAN from Springer Verlag), but their direct utility for this application seems limited due to the need for interaction with the FMRI tools. Integration with the results of the Probabilistic Atlas of the Human Brain project,<sup>16</sup> as these are made available to the community in digital format.

---

## **Spatial Normalization**

---

The normal variation in size and shape of adult human brains, coupled with the advent of high resolution quantitative imaging methods, has lead to the need for a coordinate system to describe locations in the brain. The most widely accepted such

coordinate system is adapted from the atlas of Talairach and Tournoux,<sup>9</sup> in which the rectangular box containing the cerebrum is divided into 12 subvolumes based on the anterior and posterior commissures. Each subvolume is assigned  $(x, y, z)$  dimensions based the atlas brain. The custom in neuroimaging papers has been to refer to brain images thus normalized by the millimetric coordinates of the atlas cross-sectional figures<sup>15</sup> (although this practice is deprecated in the atlas text itself).

AFNI. The transformation to Talairach coordinates is accomplished exactly as described in the atlas. The anterior and posterior commissures, and the longitudinal fissure, must be selected manually on a high resolution 3D data set. These locations define the  $(x, y, z)$  directions; the brain images are then presented aligned with these axes. The extreme points of the cerebrum (the brain bounding box) are then manually chosen. Along the with locations of the commissures, these define the 12 subvolumes, which then are separately affinely scaled to the atlas dimensions.

With high resolution (1–2 mm<sup>3</sup>) and high gray-white matter contrast images, the whole procedure takes a trained operator 2–5 minutes per anatomical data set. Lower resolution functional data sets can then be mapped to stereotaxic coordinates using the same transformation matrices. Once the transformation has been computed for a ‘parent’ anatomical dataset, `afni` will do this automatically to all other datasets (usually functional) which are linked to the parent.

The actual 3D brick arrays are not transformed until ordered by the user. For display purposes, only 2D slices interpolated at arbitrary orientation from the original 3D data are needed. This is much faster than transforming an entire set of data, and makes it possible to see slices in Talairach coordinates immediately after the requisite anatomical locations are marked.

Desiderata. Automating the transformation to Talairach coordinates has been shown to be feasible, if a database of manually transformed MRI volumes with the same contrast is available.<sup>10</sup> The technique is to maximize the correlation of the grayscale 3D brick for each subject with the average Talairach brick from the database. For a tool to be used at many sites with different scanners and pulse sequences, this may not be a sufficiently general purpose approach. One approach that might be more applicable to a wide variety of image contrasts is a preliminary probabilistic segmentation into gray and white matter, followed by a search for a transformation to bring these classifications into conformance with a ‘standard brain’.

The Talairach-Tournoux atlas does not include the cerebellum. That part of the brain is quite variable in size and shape among normal adults, due in part to its anatomical location near the base of the skull. As a result of these two factors, stereotaxic coordinates in the cerebellum brain are grossly variable between normal adults, and have little correlation with anatomy.

The Talairach coordinate system that is now used makes all 12 subvolumes of the cerebrum have the same rectilinear dimensions, which range from 23 to 79 mm. This produces an RMS deviation in the coordinate location of manually determined cortical landmarks of 6–8 mm: plus or minus one gyrus, in effect.<sup>10, 11</sup> Several efforts to map brains to a common basis at the finer sulcal-gyral level using more elaborate geometrical transformations are well underway.<sup>17, 18</sup> No consensus appears near on which approach is most generally applicable, or which one will be used for reporting brain mapping results.

Mapping the cerebral cortex onto a flat 2D surface<sup>19</sup> has some attractions both as a visualization tool and as a basis for a cortical coordinate system. Present methods require a laborious tracing of the cortical folds. This limits their widespread

application and routine use. As a stopgap and for visualization purposes, a flattening of the Talairach-Tournoux atlas brain could be used.<sup>20</sup>

There are two main issues involved in improving on the stereotaxic coordinate transformation now in use. The first issue is a matter of consensus building: what will become accepted as the baseline method for reporting brain activation maps? At a minimum, a consensus should be reached on extending the Talairach system to normalize the size and location of the cerebellum. The second issue is a matter of algorithm sharing: most work being done in this area involves intricate techniques and implementations (unlike the Talairach transformation), which are generally guarded closely. If a new system is to become a standard for the interchange of scientific information, it must be universally available.

---

## **Analysis of Functional Activation**

---

Many methods have been proposed and used for the detection of neurologically significant changes in the MR signal. Since so little of the raw signal is actually relevant, much of the effort has gone into filtering out the signals resulting from other sources.

AFNI. One major ‘other source’ is subject head motion.<sup>1</sup> A partial cure for this problem is retrospective registration of the image time series. This is accomplished in *AFNI* by the use of an iterative weighted least squares fit of each image in the sequence to a base image.<sup>5</sup> Since a large number of images is aligned to the same base image, this algorithm can be implemented very efficiently: under 50 msec per  $64 \times 64$  image on a 150 MHz Pentium system.

The correlation method for activation detection<sup>8</sup> is available interactively and in a batch program. It can be implemented very efficiently,<sup>3</sup> and can generate

3D activation maps in under a minute even from very large data sets. In the interactive mode, the correlation coefficient threshold can be adjusted interactively with immediate effects on the color overlay denoting ‘active’ voxels.

*Desiderata.* Many fMRI experiments do not truly acquire 3D data, but rather acquire 2D multislice volumes. This poses a problem for registration, since neighboring slices are usually then separated by  $\frac{1}{2}T_R$ ; even with EPI, 2–3 s apart. Collecting the data into separate 3D volumes and registering those is the approach most commonly taken, but this implicitly assumes that there is all motion occurs between the 3D acquisitions. This is a very dubious assumption for multislice imaging. Much better would be a ‘slice-into-volume’ registration method, in which each slice is treated on its own terms with respect to the 3D volume. To our knowledge, no such algorithm exists that is capable of the fine registration required for fMRI applications.

The correlation method is easily generalized to multiple correlation (*i.e.* multiple linear regression to fit time series to higher dimensional subspaces)<sup>3, 6</sup>. An implementation of a generic constrained nonlinear multivariate regression for voxel time series would be useful for many purposes, such as analysis of pharmacologically induced changes in MR signals;<sup>7</sup> however, it would be quite difficult to make such a program efficient and robust.

---

## **Reduction and Integration**

---

The statistical reduction of huge volumes of data to a manageable and comprehensible collection is a common theme in modern science, and applies with particular force to neuro-fMRI. Invention and refinement of techniques for this purpose are active areas of research in neuroimaging, where the ‘curse of dimensionality’ poses a major challenge to conventional statistical methods.<sup>13</sup>

AFNI. Several batch programs are provided for merging functional activation datasets. The simplest program just averages 3D bricks together. Another code computes the voxel-by-voxel  $t$ -tests between 2 collections of datasets, or between one collection of datasets and zero. A set of codes is also available for performing 1 way, 2 way, and 3 way analysis of variance (ANOVA) on a voxel-by-voxel basis. In addition, an experimental principal components program can compute eigen-volumes from a collection of datasets.<sup>12</sup> At this time, none of these operations can be performed interactively from the graphical interface in `afni`.

Desiderata. As FMRI data sets are integrated, an audit trail should be established so that it is possible to take a merged data set and determine exactly from what and how it was created. Statistical principles favor direct analysis from raw data to the final reduced products. This is not usually the case with current neuroimaging methods. The stages of image reconstruction and individual activation mapping are intermediate statistical procedures, so that the ANOVA tests referred to above are actually operating on already processed data. This hierarchy of processing should be traceable.

Ideally, it would be possible to examine the integrated data interactively, and at any point step backwards along the analysis chain to examine the earlier stages. At the level of regional ‘meta-analysis’, this is already possible with the program `BrainMap`.<sup>14</sup> In this system, the basic level of ‘data’ is the results from an individual scientific paper, expressed in stereotaxic coordinates as centers of activation. One type of meta-analysis is the specification of the subject stimulus conditions in order to see what regions have been reported as activated. From the results, it is possible to go backwards and determine which papers contributed.



As data and results accumulate within an institution, it will become desirable to allow the outputs of such meta-analyses—combining studies from within a unified research program—to be traced back not just to the paper level, but all the way to the image data level. Achieving this will not be easy. One challenge will be technological: designing and implementing a robust linkage system to connect the gigantic amount of data involved. Another challenge will be perceptual: understanding the need for organizing data on such a large scale.

---

## Extension and Communication

---

A true *system* comprises many parts which must be somewhat independent to be maintainable, but must also work together to be useful. In addition, the ability to build on the existing parts and integrate new capabilities is necessary for a system to have widespread research utility.

AFNI. Interactive capabilities can be added to program `afni` with the plugin facility. By writing routines in conformance with the documentation, programmers can create ‘fill-in-the-blanks’ forms to get parameters from the user, can access any of the anatomical and functional data loaded into `afni`, and can send new 3D data sets back to `afni` for display.

Desiderata. Writing a plugin for `afni` requires structuring the code in a particular way. A more general method would be to define an interprocess communication protocol for FMRI data. One possibility would be an extension of the DICOM standard;<sup>21</sup> a disadvantage of this approach is that the DICOM standard is extremely complex and lengthy (over 1000 pages). An alternative is to formalize the description of program interfaces, including methods for describing multidimensional data formats. For example, the *AFNI* package includes a facility that can read almost

any uncompressed 1D to 4D data file; it works by putting the burden on the user to specify the locations inside the file where the data is found. In this way, the program does not need to know how to decipher the many different header formats in use. The method used to describe the user interface to a plugin is similarly general. Adoption of some standards like these would allow FMRI packages from different institutions to call each other in batch mode.

A more advanced protocol would allow programs to communicate interactively. The most useful paradigm for this is ‘object oriented’ programming, and relies on the operating system to mediate communication between programs, now viewed as ‘methods’ to be applied to ‘objects’. Although this is the direction in which operating system and application design is headed, it will be some time before it is stable and platform-independent enough to be usable for a widely distributed tool that will be used in many environments.

---

## Conclusion

---

*AFNI* is a large software package, but it is perhaps only one-tenth of what is needed for a comprehensive FMRI data analysis package. Many complementary efforts are underway at other institutions and in some cases are well advanced. Cooperation among these disparate development programs could result in an widely useful and widely used system. Recent years have seen successful far-flung collaborations resulting in the development of very large and freely available pieces of software; for example, the Linux operating system and the many GNU packages.

There are several factors that are needed for such endeavors to succeed. One is the establishment of standards for communication between components. In software engineering terms, what is needed is a set of protocols to which all developers adhere; for example, the development of the HTML standard and HTTP protocol were

instrumental in recent developments on the Internet. In the context of FMRI, this means both the data file format(s) and standards for transmitting data interactively between programs.

A second requisite factor is the creation of layers of application, with defined programming interfaces. Much of what is now thought of as Unix is really X11 Windows, which is not part of Unix *per se*, but is a separate piece of software above the operating system. X11 calls upon Unix to perform certain functions (I/O and interprocess communication), and in turn provides other functions (graphics and windows) to the programmer. Layered above X11 are the X toolkit (Xt) and above that the Motif toolkit. In turn, applications such as `afni` are above all of these, and call upon the functions provided by the lower layers as needed. In the context of FMRI, this means that well planned and well documented interfaces are needed to provide a hierarchy of software functions for data access, display, processing, and auditing.

A third requisite factor is a central coordinating site, which is needed to ensure that all pieces work together and adhere to the community generated standards. This site can also function as a repository for stable versions of the software components, and for documentation. For Linux, Linus Torvalds was the coordinator; for the GNU project, the GNU Foundation and Richard Stallman served in this role.

An overriding requirement is the willingness of FMRI software development sites to cooperate. The perceived costs of cooperating are many, and the gratifications may not be instant. If the cost perceptions can be altered, then the next few years may see the creation of a very flexible and useful tools for use by the international community of FMRI researchers.

---

## Acknowledgements

---

The ongoing development of the *AFNI* package has been greatly facilitated by feedback and encouragement from many neuroscientists, including (but not limited to) MS Beauchamp, JR Binder, EA DeYoe, SM Rao, EA Stein, and FZ Yetkin. This work was partially supported by NIH grants MH51358 and NS34798.

## References

1. Hajnal JV, Myers R, Oatridge A, Schwieso JE, Young IR, and Bydder GM. Artifacts due to stimulus correlated motion in functional imaging of the brain. *Magn. Reson. Med.* **31**, 283-291 (1994).
2. Biswal B, Yetkin FZ, Haughton VM, and Hyde JS. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn. Reson. Med.* **34**, 537-541 (1995).
3. Cox RW. AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research* **29**, 162-173 (1996).
4. Cox RW, Jesmanowicz A, and Hyde JS. Real-time functional magnetic resonance imaging. *Magn. Reson. Med.* **33**, 230-236, (1995).
5. M Irani, and S Peleg. Improving image resolution by image registration. *CVGIP* **53**, 231-239 (1991).
6. Maisog JM, Courtney S, van Horn JD, and JV Haxby. Multivariate multiple regression on fMRI data to map functionally distinct areas. *NeuroImage* **3**, S79, 2<sup>nd</sup> Intl. Conf. on Functional Mapping of the Human Brain, Boston 1996..

7. Stein EA, Bloom AS, Pankiewicz J, Harsch H, Fuller SA, Cho J-K, and Rao SM. Analysis of pharmacologically-induced fMRI slow waves. *NeuroImage* **3**, S97, 2<sup>nd</sup> Intl. Conf. on Functional Mapping of the Human Brain, Boston 1996..
8. Bandettini PA, Jesmanowicz A, Wong EC, and Hyde JS. Processing strategies for time-course data sets in functional MRI of the human brain. *Magn. Reson. Med.* **30**, 161–173 (1993).
9. Talairach J, and Tournoux P. “Co-Planar Stereotaxic Atlas of the Human Brain.”, Thieme Medical Publishers, New York, 1988.
10. Collins DL, Neelin P, Peters TM, and Evans AC. Automatic 3D intersubject registration of MR volumetric data in standardized Talairach space. *J. Comput. Assist. Tomog.* **18**, 192–205 (1994).
11. Toga AW, Jones AS, Rothfeld JM, Woods RP, Payne BA, Huang C, Mazziota JC, and Cai R. Anatomic variability as measured with a 3D reconstructed Talairach atlas. In “Quantification of Brain Function, Tracer Kinetics, and Image Analysis” (K Uemura, Ed.), pp. 459–456., Elsevier Science, New York, 1993.
12. Cox RW, Binder JR, DeYoe EA, and Rao SM. Analysis of inter-trial and inter-subject repeatability in whole brain functional magnetic resonance imaging. *Proc. Soc. Magn. Reson. 3<sup>rd</sup> Mtg.*, p. 833 (Nice 1995).
13. Ford I. Some nonontological and functionally unconnected views on current issues in the analysis of PET datasets. *J. Cereb. Blood Flow Metab.* **15**, 371–377 (1995).
14. Fox PT, and Woldorff MG. Integrating human brain maps. *Current Opinion in Neurobiology* **2**, 151–156 1994.

15. Fox PT, Burton H, Raichle ME. Mapping human somatosensory cortex with positron emission tomography. *J. Neurosurgery* **67**, 34-43 (1987).
16. Thompson PM, Schwartz C, and Toga AW. High-resolution random mesh algorithms for creating a probabilistic 3D surface atlas of the human brain. *NeuroImage* **3**, 19-34 1996.
17. Collins DL, Holmes CJ, Peters TM, and Evans AC. Automatic 3-D model-based neuroanatomical segmentation. *Human Brain Mapping* **3**, 190-208 (1995).
18. Davatzikos C. Spatial normalization of 3D brain images using deformable models. *J. Comput. Assist. Tomog.* **20**, 656-665 (1996).
19. Carman GJ, Drury HA, and van Essen DC. Computational methods for reconstructing and unfolding the cerebral cortex. *Cerebral Cortex* **5**, 506-517 (1995).
20. DeYoe EA, Carman GJ, Bandettini PA, Glickman S, Wieser J, Cox RW, Miller D, and Neitz J. Mapping striate and extrastriate visual areas in human cerebral cortex. *PNAS* **93**, 2382-2386 (1996).
21. Bidgood WD, and Horii SC. Modular extension of the ACR-NEMA DICOM standard to support new diagnostic imaging modalities and services. *J. Digital Imaging* **9**, 67-77 (1996).