

tipiX: Rapid Visualization of Large Image Collections

Adrian V. Dalca¹, Ramesh Sridharan¹, Natalia Rost², and Polina Golland¹

¹ Computer Science and Artificial Intelligence Lab, MIT

² Department of Neurology, Massachusetts General Hospital, Harvard Medical School

Abstract. We present a novel approach for fast and effective visualization of large image collections in population studies. The key insight is to collapse inherently high-dimensional imaging data onto an interactive two-dimensional canvas native to a computer screen in a way that enables intuitive browsing of the image data. Increasingly, medical image computing research involves exploring large image sets with high intrinsic dimensionality. This includes three dimensions for each medical volume, and many *meta*-dimensions such as subject index, modality type in multimodal studies, time in longitudinal studies, or parameter choice in parameter sweep experiments. Current visualization tools generally display one or few 2D slices or 3D renderings at a time, and do not provide a natural way to explore the meta-dimensions. Instead, population statistics are often employed to summarize large cohorts. We propose a novel visualization approach that enables rapid interactive visualization of high dimensional image data, bridging the gap between single-volume viewers and large dataset statistics. Our approach allows users to identify important patterns in the data or anomalies that might otherwise be overlooked. We demonstrate that our platform can be used for quick and effective evaluation and analysis, and we believe it will improve research workflow and facilitate novel method development. Our tool is freely available at <http://tipix.csail.mit.edu>, where we also provide a video and live demonstration.

1 Introduction

We present a new method for rapid visualization of large medical image collections, where a two-dimensional interactive canvas is used to capture image data that is inherently high-dimensional. We aim to harness researchers' innate ability to identify visual patterns and deviations from those patterns. Specifically, our approach and the resulting visualization tool, called tipiX, bridge the gap between single-image visualization offered by most software – or *viewers* – and statistical population analysis. We achieve this goal by presenting large image collections in their entirety in an intuitive manner.

Research in medical image analysis involves increasingly large amounts of data. For example, large clinical cohorts are becoming available [6,11,14,19], offering thousands of subjects with multi-modal and longitudinal scans alongside clinical information. Parameter sweeps for processing steps can also produce large collections of images, even for a small study. These datasets include meta-dimensions – dimensions other than the three spatial ones – such as time, subject

index or parameter setting. This high dimensionality makes it challenging to identify patterns in large cohorts, to evaluate quality of processing steps, and to tune algorithm parameters.

Several powerful visualization methods and tools have been demonstrated for viewing one or few images at a time [2,3,4,10]. Visualization tools are often bundled with state of the art processing and analysis software. Some packages offer a plethora of interactive visual modules that implement powerful image analysis algorithms [10]. Others offer graphics libraries that can interact with image analysis functions [15]. Finally, some tools are built to interact with complex systems for picture archiving and communication systems (PACS) or repository system [8,13]. While powerful, all of these tools are built for visualization of few subjects at a time, and do not support large image collections.

A few recently demonstrated visualization packages have allowed visualization of several volumes at once [10,20], but users are limited to only working with as many 2D images as will reasonably tile on a computer screen. Unfortunately, such grid displays are not feasible in larger datasets with hundreds of scans.

Processing of large image collections is done instead via complex pipelines [2,12], necessitating fast evaluation of various steps. Population statistics are often used to summarize intermediate or final results, to identify outliers, and to evaluate trends [2,7,17]. However, statistics are often limited, task-specific, and do not always capture the complexity inherent in individual tasks. For example, volume measurements of structures or pathologies do not identify spatial patterns. More complex statistics can capture the spatial distribution of locations, but would not identify patterns in shape. In contrast, a researcher can often visually identify complex patterns given the appropriate visualization method.

Our approach improves on single-volume viewers by exploiting the ability of users to visually detect patterns and to identify problems across large collections. Specifically, the method provides rapid intuitive visualization of entire datasets by projecting two user-specified data dimensions onto a screen and providing convenient ways to interact with other dimensions. The tool is freely available, open-source, and does not require downloading or installation as it runs locally in modern web browsers while keeping all data and processing on the user’s computer. To the best of our knowledge, this is the first method for fast visualization of large datasets.

2 Methods

We present a visualization framework, called *tipiX*, that enables rapid interactive exploration of high dimensional image sets. Given an image collection, *tipiX* displays a two-dimensional cross-section of the data. Through simple movement of the cursor, two more dimensions can be seamlessly explored, determining the two-dimensional image that is currently displayed on the screen. As a dataset may have many meta-dimensions, the user controls which dimensions are chosen for display and navigation. In this section, we present the details of our approach and discuss implementation choices. Figure 1 provides an overview of the user interface.

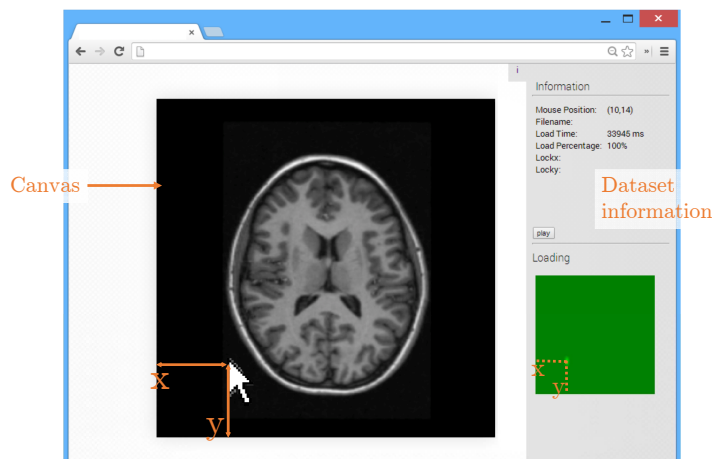


Fig. 1. User interface that implements our approach. Elements in orange were added to this screenshot for illustration purposes. The position of the cursor, shown as x and y , controls which 2D image is displayed on the canvas. The load matrix on the right information bar (green) also offers a visual indication of the position of the current image in the entire dataset.

2.1 Display Dimensions versus Navigation Dimensions

The main display – or *canvas* – shows two dimensions of the image set that users are familiar with, such as two spatial dimensions of an axial slice from a brain MRI volume. The user can simultaneously explore two more dimensions, which could be either physical dimensions or meta-dimensions. Examples include depth (the third spatial dimension), time (e.g., patient age or time of the scan), subject index in the collection, etc. For the remainder of this section, we use a dataset of 3D images with an extra dimension of subject index to illustrate the key features of the visualization method.

We use the position of the cursor on the drawing canvas to determine the location along the navigation dimensions. For example, moving the cursor to location $(4, 12)$ on the canvas displays the 4th axial slice of the 12th subject. This mode facilitates exploring two other dimensions, such as depth and time, together (if, for example, one of the navigation dimensions is set to subject age or time of the scan). The user can easily select and change display and navigation dimensions.

We avoid using sliders since they are limited to controlling one dimension at a time. This would result in much slower and more cumbersome simultaneous exploration of multiple dimensions in the datasets. In contrast, our method enables data visualization in a way that makes it easy to navigate in its entirety.

2.2 Additional Features

Our approach enables flexible selection of dimensions to be displayed. With a simple command the user can switch from viewing axial slices through volume depth and different subjects to viewing sagittal slices in the same data, as illustrated in Figure 2.




Fig. 2. Example uses of tipiX on a dataset of 3D volumes with many subjects. Left: one axial slice of a subject is shown to the user at a time. The vertical position of the cursor controls axial slice location in the same subject, and the horizontal position of the cursor controls the subject index in the study. Right: exploring 2D sagittal views in the same dataset.

Within each dimension, tipiX provides fine control through keyboard shortcuts. For example, after identifying an outlier image among 100 subjects, a user might want to explore this subject along another dimension, such as a different modality. Our framework allows *locking* the subject (or generally, the current index for any dimension) with keyboard shortcuts to enable the user to explore the other dimension for that particular subject. Once this task is completed, the user can *unlock* the subject, and continue exploring the dataset.

Our implementation includes several other useful features such as an information panel and a preview image to summarize the entire dataset. More features are possible, and we hope the tool will grow organically as different users contribute by suggesting or implementing functionality to facilitate their research. For example, seamlessly exploring a fifth dimension with the mouse wheel is a topic of future work.

2.3 Implementation

We provide a freely available and open source implementation of tipiX. The implementation runs in modern web browsers. This design decision avoids limitations associated with cross-platform functionality and software dependencies. All processing and visualization is performed client-side on the user’s computer, thus avoiding security and privacy concerns. Sharing of visualization scenes, for example between technical researchers and clinicians, is simple and reduces to simply providing a unique web address. This functionality assumes both parties have access to the same data, for example if the data is accessible online or through a server on a private network.

The interactive canvas controls and data can also be embedded into another webpage, which is useful when developing tutorials, discussing medical datasets or previewing public data. tipiX employs the XIO library [5] to support input in popular medical imaging formats, such as DICOM [9] and NIFTI [1], in addition to regular images, such as PNG and JPG. The data is rescaled to fit on the canvas, whose size adapts to each user’s browser window.

In the next section, we present several user studies and examples that evaluate the ability of users to perform certain tasks with tipiX.




Fig. 3. Outlier detection. Left panel: a typical brain scan after registration to an atlas and skullstripping, a poorly registered subject, and a simulated hypointense lesion on a subject (yellow arrow). Right plots: boxplots of time taken to identify two outliers in mis-registration and lesion detection tasks.

3 Evaluations

To the best of our knowledge, ours is the first visualization method capable of exploring large medical image collections visually. Our approach directly addresses the problem of visualizing high-dimensional data, while other viewers concentrate on specific volumes. We therefore avoid a direct comparison, as it would be unfair to the baseline methods. Instead, we demonstrate the utility of our platform through several user studies of typical medical image computing tasks. Each study included 9 users who were familiar with images, and who were provided with a brief, one-minute instruction on how to use *tipiX*. As the strengths of a visualization tool are better demonstrated visually, we encourage the readers to use the tool or view the demonstration videos available at <http://tipix.csail.mit.edu>.

3.1 Outlier Detection

As part of a typical medical image computing workflow, various processing steps must be evaluated for correctness. For example, following registration of a subject cohort, it is important to identify scans that did not register properly to an atlas template. Similarly, in a general cohort of patients, it is useful to quickly identify patients with a particular pathology.

In the first user study, we register 20 T1-weighted brain MRI scans that are part of the Freesurfer brain atlas [2] to a common atlas. We introduce a random perturbation in one registration to mimic misalignment. Users are given the entire dataset, and asked to identify a mis-registered subject in the cohort, where the mis-registration is only noticeable in some of the slices. In a similar fashion, we simulate a hypointense lesion in the white matter of one subject, and ask the users to identify the scan that contains the lesion. The lesion is only noticeable in about 11 of 100 brain slices.

These would be difficult tasks for single-volume visualization tools. Since the problems are only noticeable in a subset of the slices, re-factoring the dataset




Fig. 4. Detection of inclusion threshold. Left panel: a typical well-aligned subject image shown next to a typical poorly registered image, with the template brain and ventricle boundaries shown in green. Right plot: the average threshold versus the time taken by users.

to a single volume of a specific slice across subjects is not feasible. Figure 3 illustrates example volume slices and reports the user time required for outlier detection.

We find that users are very adept at identifying outliers when using our visualization tool, taking on average about 45 seconds to identify a mis-registration, and about 90 seconds to identify a subject with a lesion. We observed empirically that users who first explored the dataset across volumes rather than within a single volume tended to complete the tasks faster. This illustrates the power of our approach of visualizing meta-dimensions and image dimensions simultaneously. Quick visual detection promises to substantially improve quality control in complex tasks, where designing robust quality control measures can be time consuming or prohibitive.

3.2 Pattern Identification

Similarly to identifying outliers following a pipeline step, we often aim to subdivide a subject cohort. For example, after rigid registration of a large image set to an atlas, we may wish to identify a subset of subjects that have promising registration results for analysis. In this study, we rank 500 non-rigidly registered subjects from a large neuroimaging study of stroke [16] based on the sum of squared differences between each deformed subject image and the atlas template. Using this ranking, our method allows users to explore the entire dataset and identify a threshold on the similarity measure that would separate well-registered subjects from those with significant misalignment. In Figure 4, we illustrate typical good and poor registrations and report the agreement level among different users in the study. We found that nearly all users decided quickly on a rough threshold range, but some spent more time deciding on an optimal choice and remarked on the subjective nature of the task.

Parameter exploration is necessary in many applications, such as processing of new datasets or development of novel methods. Registration quality metrics




Fig. 5. A screenshot of two images shown to users in the parameter sweep study. Axial slices from four separate subjects are shown with segmentation contours (orange) implied by the aligned atlas. Moving the cursor left to right sweeps the registration parameters; moving the cursor up and down goes through different axial slices. Left: results from the median parameter choice. Right: a parameter setting chosen by a user who noted that the top-left subject is otherwise mis-registered.

often do not capture complex patterns or task-specific problems. Visualization of the behaviour of various parameters on a subset of the population can help to quickly identify the optimal value of the regularization parameter choice. We perform a parameter sweep for non-rigid registration using the Log Domain Diffeomorphic Demons algorithm [18] on four subjects, and ask users to identify the optimal parameter (Figure 5). We found that most users chose very similar settings (smoothing kernel size near 1.5mm) in about two minutes, and spent the majority of that time refining the optimal parameter once a reasonable range was found.

4 Conclusion

In this paper, we demonstrated a novel approach to visualizing high-dimensional medical image sets. Our method enables visualization of spatial dimensions and *meta*-dimensions, including different patients and modalities. Moreover, analysis methods often generate new dimensions that need to be explored and understood, such as growth patterns of pathologies segmented by a new algorithm. By translating the position of a cursor on a canvas to an index into relevant image or meta-dimensions, we allow interaction and harness the human ability to quickly detect patterns and deviations from the patterns in images. We have demonstrated the usefulness of our method on several common medical imaging tasks, where we showed that users can quickly identify outliers and choose appropriate parameters from our visualizations.

Acknowledgements. This research was supported in part by: NSERC CGS-D, NSF GRFP, NIH NIBIB NAC P41EB015902, and NIH NIBIB NAMIC U54-EB005149, NIH NINDS NS082285, NIH NINDS K23 NS064052, NIH NINDS U01NS069208 and the American Stroke Association-Bugher Foundation Centers for Stroke Prevention Research. We also thank the tipiX users for participating in our studies.

References

1. RW Cox, J Ashburner, H Breman, K Fissell, C Haselgrove, et al. A (sort of) new image data format standard: Nifti-1. *Hum. Brain Mapp.*, 25, 2004.
2. B Fischl. Freesurfer. *Neuroimage*, 62(2):774–781, 2012.
3. KJ Friston, JT Ashburner, SJ Kiebel, TE Nichols, and WD Penny. *Statistical Parametric Mapping: The Analysis of Functional Brain Images*. Acad. Press, 2011.
4. D Haehn. Slice:drop: collaborative medical imaging in the browser. In *ACM SIGGRAPH 2013 Computer Animation Festival*, pages 1–1. ACM, 2013.
5. D Haehn, N Rannou, B Ahtam, P Grant, and R Pienaar. Neuroimaging in the browser using the X toolkit. *Neuroinformatics*, 2012.
6. HB Hubert, M Feinleib, PM McNamara, and WP Castelli. Obesity as an independent risk factor for cardiovascular disease: a 26-year follow-up of participants in the Framingham heart study. *Circulation*, 67(5):968–977, 1983.
7. Z Liu, Y Wang, G Gerig, S Gouttard, R Tao, et al. Quality control of diffusion weighted images. In *SPIE Medical Imaging*, pages 76280J–76280J. International Society for Optics and Photonics, 2010.
8. DS Marcus, TR Olsen, M Ramaratnam, and RL Buckner. The extensible neuroimaging archive toolkit. *Neuroinformatics*, 5(1):11–33, 2007.
9. NEMA, ACR, et al. *Digital imaging and communications in medicine (DICOM)*. National Electrical Manufacturers Association, 1998.
10. S Pieper, M Halle, and R Kikinis. 3D Slicer. In *Proc IEEE Int Symp Biomed Imaging*, pages 632–635. IEEE, 2004.
11. EA Regan, JE Hokanson, JR Murphy, B Make, DA Lynch, et al. Genetic epidemiology of COPD (COPDgene) study design. *COPD: Journal of Chronic Obstructive Pulmonary Disease*, 7(1):32–43, 2011.
12. DE Rex, JQ Ma, and AW Toga. The LONI pipeline processing environment. *Neuroimage*, 19(3):1033–1048, 2003.
13. A Rosset, L Spadola, and O Ratib. OsiriX: an open-source software for navigating in multidimensional DICOM images. *Journal of Digital Imaging*, 17(3):205–216, 2004.
14. NS Rost, K Fitzpatrick, A Biffi, A Kanakis, W Devan, et al. White matter hyperintensity burden and susceptibility to cerebral ischemia. *Stroke*, 41(12):2807–2811, 2010.
15. WJ Schroeder, LS Avila, and W Hoffman. Visualizing with VTK: a tutorial. *IEEE Comput. Graph. Appl.*, 20(5):20–27, 2000.
16. R. Sridharan, A.V. Dalca, K.M. Fitzpatrick, L. Cloonan, A Kanakis, O. Wu, K.L. Furie, J. Rosand, N.S. Rost, and P. Golland. Quantification and analysis of large multimodal clinical image studies: Application to stroke. In L. Shen, T. Liu, P.T. Yap, H. Huang, D. Shen, and C.F. Westin, editors, *Multimodal Brain Image Analysis, LNCS*, volume 8159, pages 18–30. Springer, 2013.
17. T Stöcker, F Schneider, M Klein, U Habel, T Kellermann, et al. Automated quality assurance routines for fMRI data applied to a multicenter study. *Hum. Brain Mapp.*, 25(2):237–246, 2005.
18. T Vercauteren, X Pennec, A Perchant, and N Ayache. Symmetric log-domain diffeomorphic registration: A demons-based approach. In *Medical Image Computing and Computer-Assisted Intervention*, pages 754–761. Springer, 2008.
19. MW Weiner, DP Veitch, PS Aisen, LA Beckett, NJ Cairns, et al. The Alzheimer’s Disease Neuroimaging Initiative: A review of papers published since its inception. *Alzheimer’s & Dementia*, 9(5):e111–e194, 2013.
20. G Yang. MRI Watcher, <http://www.nitrc.org/projects/mriwatcher/>.