

# Force Brushes: Progressive Data-Driven Haptic Selection and Filtering for Multi-Variate Flow Visualizations

B. Jackson<sup>†</sup> and D. Coffey<sup>‡</sup> and D. F. Keefe<sup>§</sup>

University of Minnesota, Department of Computer Science and Engineering

---

## Abstract

We present *Force Brushes*, a haptic-based interaction technique for controlled selection in multi-variate flow visualizations. *Force Brushes* addresses the difficult task of volumetric selection and filtering by rendering haptic constraints that allow scientists to snap directly to proxy geometry, such as streamlines, to select regions of interest and then progressively filter the selection using a data-driven approach. Using progressive brushing actions with multiple variables, a user has the potential to explore volumetric data in a more immediate, fluid, and controllable way guided by the underlying data.

---

## 1. Introduction

Many scientific visualization applications rely upon user inputs to perform complex spatial tasks, such as selecting subsets of data or volumes of interest. This type of user interaction can play a critical role in the success of the visualization, enabling users to query, explore, and call up detailed data on demand. When working with volumetric datasets, traditional desktop, mouse, and keyboard interfaces can sometimes be used for these tasks; however, research has shown that traditional 2D interfaces can be very difficult to use for real-world scientific tasks, such as selecting bundles of neural fiber tracks in dense 3D DT-MRI visualizations [ZSZ\*06]. Our research explores the potential of novel haptic 3D user interface techniques for improving common interactive visualization tasks, such as volumetric filtering and selection, by making them more immediate, fluid, and controllable.

Our work addresses two problems with current 3D user interfaces for volumetric selection and filtering. The first is control. Although using 3D input to interact with 3D datasets seems intuitive, muscular jitter, fatigue, and tracking error often make 3D interfaces difficult to control. This is a major problem for working with scientific data, where precise selections are required. Our approach uses haptic feedback to improve the accuracy of both an initial 3D selection and pro-

gressive refinements to that selection. The second problem with current interfaces is a lack of a connection to the underlying data. Most of the existing 3D user interfaces that can be applied to visualization tasks take the approach of using direct input from the hand(s) to place boxes (e.g., [SAM\*05]), a lasso (e.g., [KZL08]), or some other widget in the 3D space of the visualization, then rely only upon the user's visual perception and motor control to precisely position these widgets. Our approach combines 3D user input with constraints defined interactively by multiple underlying data variables in order to enable selections of volumes that match the underlying spatial characteristics of the data.

We call the specific scientific visualization user interface technique introduced in this paper *Force Brushes* (Figure 1). Brushing – a user interface technique that involves moving a user-controlled cursor over data values to highlight or select them – has been applied previously in both 2D (e.g., [DGH03]) and 3D (e.g., [HWR\*09]) visualizations and is an intuitive strategy for selecting a subset of data. Our approach extends this concept in several ways. We contribute a 3D brushing interface for data visualization that combines the following techniques: (1) using force feedback defined by features in the underlying dataset to guide the user in making an initial feature selection, (2) using force feedback to help control a 3D pulling gesture that smartly expands the selection to include regions with similar data characteristics, and (3) using a series of these brushing operations performed in sequence with different data variables to progressively refine the selection and explore the dataset.

---

<sup>†</sup> bjackson@cs.umn.edu

<sup>‡</sup> coffey@cs.umn.edu

<sup>§</sup> keefe@cs.umn.edu

Although the resulting interface technique could eventually be applied to many different datasets, we have focused within the work presented in this paper on the common case of 3D flows represented by multi-variate vector fields, such as the hurricane dataset shown in Figure 1. To provide an organizing visual and geometric structure for these data that is larger than a voxel, force brushes operate on integral curves (e.g., streamlines) through the data, with multiple data variables sampled along these curves. In other datasets, additional basic geometric data features may also be useful (e.g., vortex core lines, fibers in muscular and other tissue visualizations, neural fiber tracts in brain visualizations).

We begin the remainder of the paper with a discussion of related work in haptic-aided visualization and volumetric selection interfaces. Then, we describe how force brushes can be used to first precisely select individual features, and then, using multiple brushes each tied to a different underlying data variable, smartly extend the selection to include volumes of similar data. Finally, we close with a discussion of opportunities for future work.


## 2. Related Work

### 2.1. Haptic Interaction for Scientific Visualization

Haptic force-feedback has previously been used as a technique for interacting with medical, art, and scientific visualizations. For example, the nanomanipulator system [TRC\*93] allowed scientists to feel microscopic surface details examined through a Scanning Tunneling Microscope. Similarly, [AS96] used haptic devices to feel the shape of a neuron model produced with a confocal scan. In contrast to these techniques, our approach enables scientists to touch abstract data such as a vector field within a flow volume, which does not have an explicit surface.

Here, we build upon previous work (e.g. [vRBKB03, IBHJ03, LSCY05, MAB08, AMBZ\*11]) that used haptic constraints to force the input pen to follow integrated streamlines through the flow vector field. Our use of haptics shares similar motivation to that of [LSCY05], which allowed users to trace the path of a vortex in CFD data. The haptic feedback complements the visual representation of the data, while not suffering from problems of occlusion and clutter. However, rather than simply feeling the flow direction, our interface leverages the haptic forces to provide more controlled selection of regions of interest, as well as to provide input to control the visualization.

This approach shares similar motivation to the springs and constraints for 3D drawing by Snibbe et al. [SSV] and Drawing on Air [KZL07], which used non-realistic forces to allow for more controlled input. However, our application area targets the exploration of scientific data, rather than 3D artistic creation.



**Figure 1:** Force Brushes are used to explore a simulated multi-variate hurricane dataset. The user works in front of a 3D visualization with a SensAble Phantom Premium 1.5 device, holding the stylus in his dominant hand. Key presses made on a small keypad by the non-dominant hand are used to select different brushes, with each brush tied to a specific variable in the underlying dataset.

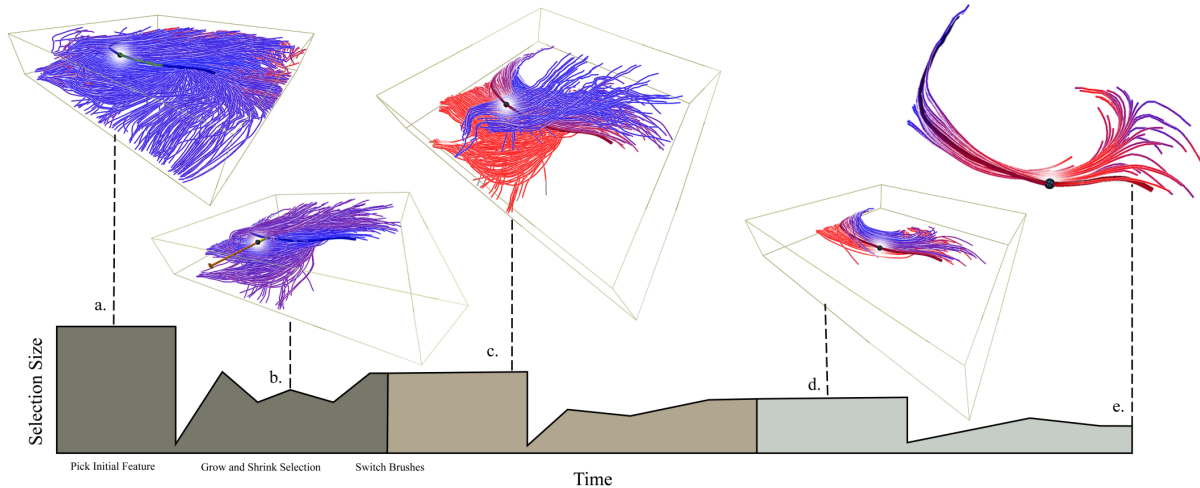
### 2.2. Streamline Selection

Several works, such as [SFL\*04, SAM\*05, Ake06, WZL08, CML\*11], have examined the task of selecting streamlines in flow visualizations or dense DTI fiber tracts. This remains an active area of research. Predominately, these relied on spatially selecting volumes of interest with lassos, boxes or other widgets. Of particular note, is the haptic-assisted 3D lasso drawing system by Zhou et al. [WZL08]. This system made use of a force-feedback device to provide precise curve drawing for 3D lassos; however, in contrast to our approach it did not use the underlying data to help constrain the selection, relying only on the scientist drawing free-form curves.

Most similar to our approach of creating a selection based on a specific streamline in the data is the shape marking operation in the CINCH system [Ake06]. This operation allows the user to draw a representative 2D line on a tablet and returns a selection of similar lines. Although our technique of selecting a representative line from the data is more constrained than drawing a free-hand line, we avoid any ambiguity associated with projecting a drawn 2D line into a 3D volume. In addition, we enable progressive refinement of the simulation based on multiple variables and show how force-feedback can be used in this context.

## 3. Visualizing Hurricane Isabel: A Use Case For Multi-Variate Flow Feature Selection

An ideal use case for multi-variate flow selection is visualizing and analyzing weather data. These data frequently contain multiple variables such as wind speed, pressure, amount of moisture, and temperature. Our application, shown in Fig-



**Figure 2:** A sample workflow using Force Brushes. From left to right the selection is refined using a series of three brushes. For each brush, the user first selects an initial feature of interest, either an entire line or a subset of a line as in (a); then pulls out away from the line to grow or shrink the selection as in (b); then releases the stylus button to push the current selection to the stack as in (c). Additional brushes can be applied to further refine the selection based on other data variables as in (d) and (e).

ure 1, allows scientists to selectively filter the visualization to identify critical regions of interest.

To do this, the typical workflow involves first picking a brush, such as *similar shape* or *average wind speed*. This updates the visualization to display just these data. Then, the haptic stylus is moved into the flow volume to explore the data and find a feature of interest. A magic lens clips the geometry between the camera and the stylus enabling the user to see deep into the flow volume. The user picks a feature of interest by pressing and holding the stylus button. To indicate a subregion of a line, the user drags along the line, as in Figure 2a. Next, to perform a selection based on the identified feature, the user pulls the brush away from the line, as in Figure 2b. This has the effect of selecting only the single initial feature of interest identified by the user (e.g., one line), but as the user pulls the stylus further away from the initial line, the selection grows to include data from similar lines, where similarity is defined by the underlying data and the brush type (e.g., pressure, velocity, temperature). This process is fluid; in order to pick just the right subset of data, the user will typically repeatedly pull the brush far away and then move it back closer to the line while watching the selection update dynamically, as shown in the accompanying video. To complete the operation, the user releases the stylus button, which pushes the active selection onto a stack (Figure 2c). Then, using the same interface, the user can continue to refine the selection with additional brushes. For example, in Figure 2, points (c), (d), and (e) each show the end result of applying a brush to the data. When applied in succession, these brushes can be used to precisely refine a selection based on several data variables. With this typical high-level

workflow in mind, the following section describes in detail the force feedback, data-driven constraints, and other technical concepts needed to make Force Brushes work.


## 4. Force Brushes

We describe in detail the two sequential tasks performed for each brush: selecting a region of interest along a streamline to filter the visualization, and growing the selection. To make switching between brushes quick, each brush is mapped to a separate number on a wireless number pad (Figure 1). While one hand controls the haptic pen, the other is free to swap instantly between different brushes with a simple tap.


### 4.1. Selecting an Initial Feature

After snapping to a streamline to select it, each type of force brush can be used to select a region of interest. By pressing the button on the haptic pen, the scientist brushes along the line, releasing the button to end the selection length (shown in Figure 3). Haptic constraints keep the pen anchored to the selected streamline. In our implementation this is accomplished using an OpenHaptics surface constraint with a snap distance of 20.0. After the selection is completed, a metric for the selected feature is calculated based on the brush type. For instance, if the average pressure brush was used, the average pressure along the length of the selection is calculated, while the peak temperature brush calculates the maximum temperature in the selected range. The group of lines that are shown is then filtered to only contain streamlines that contain the same calculated value as the selected region.

Selection based on shape works slightly differently than



**Figure 3:** Dragging the haptic brush along a streamline selects a subset of the line, shown in green.



**Figure 4:** Pulling perpendicularly away from the initially selected feature grows the selection volume.

the other brushes, which operate on one dimensional values. There are several previously published methods to calculate similarity between lines (e.g. [CGG04,BPKW03,DGA01]). We use a similarity metric that was originally designed to compare DTI fiber tracts [DL09], an application area that is also applicable to our technique. This measure is based on a weighted average of distances from each sample point along the selected streamline to the closest point on another, in order to determine how closely two streamlines follow a similar path. The distances near the end points of the line are weighted more heavily. To speed up computations involving comparison between entire streamlines, we pre-compute a similarity matrix that includes the similarity between each streamline and every other.

#### 4.2. Growing the Selection

The selection growing gesture is distinguished from initial feature selection by pulling perpendicular to the selected streamline rather than along it (Shown in Figure 4). Once the point of the haptic pen leaves the line, it is constrained to a plane perpendicular to the streamline. The distance between the pen tip and the point where the streamline intersects the plane is normalized to  $(0, 1]$  using a maximum distance of 40 percent of the haptic workspace. With our Phantom Pre-

mium 1.5 haptic device this presents a total movement of about 16cm. The normalized distance is used to add streamlines back to the selection. For example, if the currently enabled brush is average pressure, the average pressure for each streamline is calculated. The range of streamline average pressures is mapped to the normalized distance, so that at a value of one all the lines would be added to the selection, while a value closer to zero would only add lines that have average pressures very similar to the selected region.

#### 4.3. Providing Visual Feedback

To help guide the user, we color map the streamlines based on their similarity. There are two choices for this mapping: absolute coloring relative to all of the data, or relative coloring based on similar features to the currently selected line. For some variables, such as shape, only relative color mapping makes sense. Our implementation uses relative coloring for shape and absolute coloring for the other variables.

To help with streamline selection we render a magic lens around the point of the haptic pen. This lens uses a gaussian opacity filter to allow the user to continue to see the selected line if it would otherwise be occluded by lines in front of it. See Figure 3. During the growing operation, we render a ruler (shown in orange in Figure 4) extending perpendicularly to the selected line to provide spatial clues for how far the haptic pen tip (shown as a sphere) has been pulled away from the selected line.

#### 5. Conclusions and Future Work

We have presented Force Brushes, a haptic-aided technique for 3D selection and interactive data exploration in flow volume visualizations. Force Brushes are motivated by a need for more precise control over selection and other interactive data querying operations required by 3D visualizations. Although we have yet to formally evaluate the precision and control provided by Force Brushes as compared to other related approaches, our experience with the interface suggests that the combination of data-driven haptic constraints, haptic-assisted growing of selections, and progressive applications of multiple data brushes enables fast and accurate selections in dense 3D visualizations. In the future, we would like to verify this initial assessment through case study applications that range from flow visualizations to neural fiber tract visualizations and other medical applications where current 3D data querying interfaces are often a bottleneck.

#### 6. Acknowledgments

We thank Prince Bajracharya for his help implementing the streamline shape similarity metric and insights from developing an early prototype of the selection interface. Hurricane Isabel data produced by the Weather Research and Forecast (WRF) model, courtesy of NCAR and the U.S. National Science Foundation (NSF).

## References

- [Ake06] AKERS D.: CINCH: a cooperatively designed marking interface for 3D pathway selection. In *Proceedings of the 19th annual ACM symposium on User interface software and technology* (Montreux, Switzerland, 2006), pp. 33–42.
- [AMBZ\*11] ABDUL-MASSIH M., BENEŠ B., ZHANG T., PLATZER C., LEAVENWORTH W., ZHUO H., GARCÍA E. R., LIANG Z.: Augmenting heteronanostructure visualization with haptic feedback. In *Proceedings of the 7th International Conference on Advances in Visual Computing - Volume Part II* (2011), pp. 627–636.
- [AS96] AVILA R. S., SOBIERAJSKI L. M.: A haptic interaction method for volume visualization. In *Proceedings of the 7th conference on Visualization '96* (1996), VIS '96, pp. 197–ff.
- [BPKW03] BRUN A., PARK H.-J., KNUTSSON H., WESTIN C.-F.: Coloring of DT-MRI fiber traces using laplacian eigenmaps. In *Computer Aided Systems Theory - EUROCAST 2003* (2003), vol. 2809, pp. 518–529.
- [CGG04] COROUGE I., GOYTARD S., GERIG G.: Towards a shape model of white matter fiber bundles using diffusion tensor MRI. In *Biomedical Imaging: Nano to Macro, 2004. IEEE International Symposium on* (2004), pp. 344–347.
- [CML\*11] COFFEY D., MALBRAATEN N., LE T., BORAZJANI I., SOTIROPOULOS F., ERDMAN A., KEEFE D.: Interactive Slice WIM: Navigating and interrogating volume datasets using a multi-surface, multi-touch VR interface. *IEEE Transactions on Visualization and Computer Graphics* PP, 99 (2011), 1.
- [DGA01] DING Z., GORE J. C., ANDERSON A. W.: Case study: reconstruction, visualization and quantification of neuronal fiber pathways. In *Proceedings of the conference on Visualization '01* (2001), VIS '01, pp. 453–456.
- [DGH03] DOLEISCH H., GASSER M., HAUSER H.: Interactive feature specification for focus+context visualization of complex simulation data. In *Proceedings of the symposium on Data visualization 2003* (2003), VISSYM '03, pp. 239–248.
- [DL09] DEMIRALP C., LAIDLAW D. H.: Similarity coloring of DTI fiber tracts. In *Proceedings of DMFC Workshop at MICCAI* (2009).
- [HWR\*09] HENTSCHEL B., WOLTER M., RENZE P., SCHRÄUDER W., BISCHOF C. H., TORSTEN K.: Hybrid parallelization for multi-view visualization of time-dependent simulation data. In *Proceedings of EGPGV'2009* (2009), pp. 79–86.
- [IBHJ03] IKITS M., BREDERSON J., HANSEN C., JOHNSON C.: A constraint-based technique for haptic volume exploration. In *IEEE Visualization* (2003), pp. 263–269.
- [KZL07] KEEFE D. F., ZELEZNIK R. C., LAIDLAW D. H.: Drawing on air: Input techniques for controlled 3D line illustration. *IEEE Transactions on Visualization and Computer Graphics* 13, 5 (2007), 1067–1081.
- [KZL08] KEEFE D. F., ZELEZNIK R. C., LAIDLAW D. H.: Tech-note: Dynamic dragging for input of 3D trajectories. In *Proceedings of IEEE Symposium on 3D User Interfaces 2008* (2008), pp. 51–54.
- [LSCY05] LUNDIN K. E., SILLEN M., COOPER M. D., YNNERMAN A.: Haptic visualization of computational fluid dynamics data using reactive forces. In *Proceedings of the Conference on Visualization and Data Analysis, part of IS&T/SPIE Symposium on Electronic Imaging* (2005), vol. 5669.
- [MAB08] MÉNÉLAS B., AMMI M., BOURDOT P.: A flexible method for haptic rendering of isosurface from volumetric data. In *Proceedings of the 6th International Conference on Haptics: Perception, Devices and Scenarios* (2008), EuroHaptics '08, pp. 687–693.
- [SAM\*05] SHERBONDY A., AKERS D., MACKENZIE R., DOUGHERTY R., WANDELL B.: Exploring connectivity of the brain's white matter with dynamic queries. *Visualization and Computer Graphics, IEEE Transactions on* 11, 4 (2005), 419–430.
- [SFL\*04] SOBEL J. S., FORSBERG A. S., LAIDLAW D. H., ZELEZNIK R. C., KEEFE D. F., PIVKIN I., KARNIADAKIS G. E., RICHARDSON P., SWARTZ S.: Particle flurries: synoptic 3D pulsatile flow visualization. *IEEE Computer Graphics and Applications* 24, 2 (2004), 76–85.
- [SSV] S. SNIBBE S. A., VERPLANKI B.: *Springs and Constraints for 3D Drawing*. Tech. rep.
- [TRC\*93] TAYLOR R. M., ROBINETT W., CHI V. L., BROOKS JR. F. P., WRIGHT W. V., WILLIAMS R. S., SNYDER E. J.: The nanomanipulator: a virtual-reality interface for a scanning tunneling microscope. In *Proceedings of the 20th annual Conference on Computer Graphics and Interactive Techniques* (1993), SIGGRAPH '93, pp. 127–134.
- [VRBKB03] VAN REIMERSDAHL T., BLEY F., KUHLEN T., BISCHOF C. H.: Haptic rendering techniques for the interactive exploration of CFD datasets in virtual environments. In *Proceedings of the Workshop on Virtual Environments 2003* (2003), EGVE '03, pp. 241–246.
- [WZL08] WENJIN ZHOU S. C., LAIDLAW D. H.: Haptics-Assisted 3D lasso drawing for tracts-of-interest selection in DTI visualization. In *IEEE Visualization Poster Compendium* (2008).
- [ZSZ\*06] ZHOU W., SIBLEY P. G., ZHANG S., TATE D., LAIDLAW D. H.: Perceptual coloring and 2D sketching for segmentation of neural pathways. *Siggraph 2006, Poster Compendium*, 2006.