Kinetica: Naturalistic Multi-touch Data Visualization

Jeffrey M. Rzeszotarski, Aniket Kittur
Human-Computer Interaction Institute, Carnegie Mellon University, Pittsburgh, PA, USA
{jeffrz, nkittur}@cs.cmu.edu

ABSTRACT
Over the last several years there has been an explosion of powerful, affordable, multi-touch devices. This provides an outstanding opportunity for novel data visualization techniques that leverage new interaction methods and minimize their barriers to entry. In this paper we describe an approach for multivariate data visualization that uses physics-based affordances that are easy to intuit, constraints that are easy to apply and visualize, and a consistent view as data is manipulated in order to promote data exploration and interrogation. We provide a framework for exploring this problem space, and an example proof of concept system called Kinetica. We describe the results of a user study that suggest users of Kinetica were able to explore multiple dimensions of data at once, identify outliers, and discover trends with minimal training.

Author Keywords
Visualization, multivariate data, multi-touch, physics

ACM Classification Keywords
H.5.m. Information interfaces and presentation: Misc.

INTRODUCTION
The ways in which people consume information and make decisions is changing rapidly, with a rise in powerful, affordable, multi-touch devices that support a wide range of interactions and applications. A 2013 Pew Research survey estimated that one third of American households owned a tablet computer [40], a number double that of the previous year. There has also come a bevy of collaborative surfaces, interactive whiteboards, and touch-enabled traditional PCs. These devices create an opportunity for developing natural user interfaces (NUIs) that are grounded in naturalistic human ways of interacting with real things [12].

However, despite the crucial importance of interpreting and understanding complex data in nearly any personal or business setting [31], existing information visualization techniques have largely failed to make effective use of the multi-touch, direct manipulation capabilities of these new, ubiquitous devices [23]. Instead, techniques remain

CHI 2014, April 26 - May 01 2014, Toronto, ON, Canada

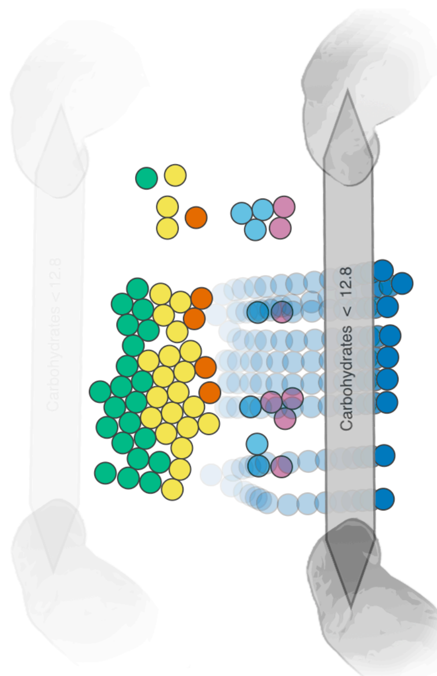

Figure 1: A user drags a two finger semi-permeable filter across some data. Points that meet its criteria collide and are pulled with it to the right, passing over unfiltered points.

primarily within the realm of traditional, directed interactions using WIMP metaphors. Lee et al. employed the terms *post-WIMP* and *post direct manipulation* to describe the rich area of research currently filling this gap [23]. Such techniques that move towards gestures and naturalistic interactions may have benefits for information literacy and awareness. For exploratory data visualization, post-WIMP approaches have the potential to preserve users' growing mental model of an information space as they explore it, and help them interact with more variables more fluidly than traditional interfaces. This might help them understand the structure and distributions of data even without significant training or statistical expertise by leveraging their models of the physical world [18].

In this paper we explore how post-WIMP interactions might improve exploratory data visualization through the introduction of physics-based affordances and multi-touch interaction techniques. Grounding data exploration in a physics-based metaphor may have a number of benefits, including supporting richer mental models of an information space by keeping data salient and fluidly tracking their locations over time; providing better awareness of amounts and distributions; making outliers stand out over the course of exploration; and augmenting working memory through visual traces. It may also minimize training time and participation barriers by tapping into users' knowledge of the physical world, their own bodies, their surroundings, and their social context [18]. We describe a proof of concept system, Kinetica, and use it to demonstrate some of the benefits of this approach through a comparative user study. We find that physics-based interactions help users explore multiple dimensions at once, make more descriptive and comparative findings about their data, and develop a more holistic understanding of a dataset. We also present a generative framework for naturalistic multi-touch visualization that may help to inspire future work in this space. Finally, we consider limitations and discuss areas for future work.

RELATED WORK

Below we briefly discuss some of the most relevant related work covering different aspects of post-WIMP information visualization and systems for exploring data.

Multi-touch and Gestural Interaction

Gestures, multi-touch, and physical interactions are key components of NUIs regardless of the particular application. Keefe and Isenberg [20] highlight some of the major interest areas for natural interaction in data visualization. Foremost are actual ways of interacting with data, which Isenberg and Isenberg [17] break down into the relationship between data representation and interaction technique. What are the interaction primitives one uses, and how do we bind them to data visualization operations? RoomPlanner [38] used physical actions like cupping a hand and transparency to make data easy to explore on a DiamondTouch surface, and Wilson et al. explored a wide variety of game physics-based interactions [35]. These interactions have the benefit of both ease of use and the ability for users to improvise new ones based on their experiences in the real world. Wobbrock et al. [36] take this idea further, allowing users to build their own gestures.

Moving towards more structured interaction techniques, NEAT and interactive grids [9,10] provide a set of pen and touch primitives to align and locate data on a field. This can be critical given the imprecision that can accompany touch interaction. TouchWave [3] displays stacked hierarchical graphs much in the way a WIMP application might, but uses drags and swipes to scale and separate overlaps naturalistically. Schmidt et al. [28] treat graph edges and nodes like beads on strings that can be gathered and stacked. SketchInsight and SketchStory [24,34] combine interactive sketching with dynamically constructed traditional charts for didactic presentations. Some approaches even employ physical objects. Senseboard [19] places pucks on a grid to organize and manipulate information, while Ullmer et al. [33] use physical tokens to control and visualize database queries.

Another crucial component Isenberg and Isenberg identify is evaluation. Drucker et al. [6] provide one such example. They use a grounded methodology to study user preference and effectiveness in different visualization prototypes in comparison, findings NUIs to be effective but at times constraining when people already knew what to do.

Multivariate Data Visualization

Wong and Bergeron provide a good summary of the history of multivariate visualization [37]. Early examples such as starfield visualizations [2], brushing [26], and dynamic querying [30] permit similar interactions, allowing users to filter different dimensions with active feedback. To visualize complex or categorical data, techniques such as parallel coordinates [16] and parallel sets [22] offload multivariate relationships onto a visual field, highlighting outliers and trends. DataMeadows [7] and Jigsaw [32] all use multiple techniques in tandem to represent trends or clusters within multiple dimensions.

One highly relevant inspiration for our line of work is that of Yi et al.'s Dust & Magnet [39], which used a magnet physical metaphor to visualize multivariate data. By allowing users to place and move magnets, attracting points, users could discover trends among their data. We build on the idea of using physics and interaction to compress the dimensionality of multivariate data, but generalize to a wider set of physics-based techniques.

Physics and Interactivity

One avenue for intuitive, interesting interactions cited by Jacob et al. [18] is the introduction of physical metaphors. Bumptop [1] shows one such example in a general computing context. Files become physics objects that can be pinned to walls and moved around in an environment. This has the benefit of requiring minimal training, since people already intuit that objects fall and pins stick things to walls. North et al. [27] extend this sort of physics approach into data sorting in comparison with a traditional WIMP system. People readily adapted to physical models, outperforming those using a mouse. Force-based graph layouts use attraction and repulsion, which people naturally intuit, to make graph exploration easier [11]. Sticky tools make touch manipulations of objects more naturalistic and therefore easier [13]. Visual sedimentation metaphors use the buildup of material to analogize the aggregation of time-series data in an easily comprehensible way [15].

Klemmer, Hartmann, and Takayama [21] suggest a more general explanation of this sort of effect. The *tangibility* of an artifact, or how naturally it maps actions to physical reality can greatly influence the performance of a system. While closer mappings to reality may come with additional constraints and limitations, they also provide benefits for training and understanding, especially for inexpert users. This might extend beyond just NUIs. The description of post-WIMP possibilities that Jacobs et al. provide and the design considerations of Lee et al. [23] still might work within traditional mouse-and-keyboard computers. While they may not be quite as *tangible* as if one were touching

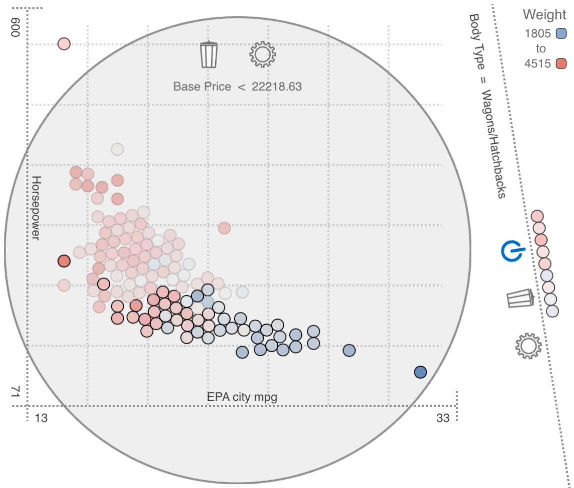

Figure 2: A user explores car model data using Kinetica. She cares about mileage and power, so she placed the points into a scatterplot. She doesn't like wagons, so she used a wall to filter. Finally, she added a lens to highlight cheaper vehicles.

the screen, the interactions and feedback may nonetheless remain physically grounded and beneficial.

KINETICA

To explore the benefits of physics-based affordances and multi-touch interactions for data visualization and to understand the technical challenges of implementing such techniques, we developed a proof of concept application called Kinetica. For our proof of concept, we used a tablet computer because it occupied the physical space of the user, could be twisted and turned, and was responsive to touch. We decided to use an iPad due to the convenience and general availability of the device, though we hope to extend Kinetica to more systems in the future.

We implemented Kinetica using Objective C, and make use of an MIT licensed physics library called Chipmunk Physics [25] to handle forces and collisions. This library makes use of multiple CPU cores to improve efficiency. Indeed, the most computationally costly portion of the prototype was actually the rendering loop.

We adopted an iterative strategy in developing the physics-based interactions in Kinetica. Initially we encoded data into circles colliding in a sandbox, implemented a touch-responsive magnet tool much like Dust & Magnet [39], and added gravity based on device orientation. This prototype by itself was evocative. By tilting the device so gravity took hold and pulling points with a magnet such that the forces balanced each other, data readily sorted itself and separated, highlighting outliers. We embarked in a process of loading different canonical multivariate datasets into the program and exploring them. We used datasets of cereal and car brands to explore comparisons and decision-making. We used datasets of Titanic shipwreck passengers and Pittsburgh census demographics to explore how one might generate and test hypotheses about trends in the data.

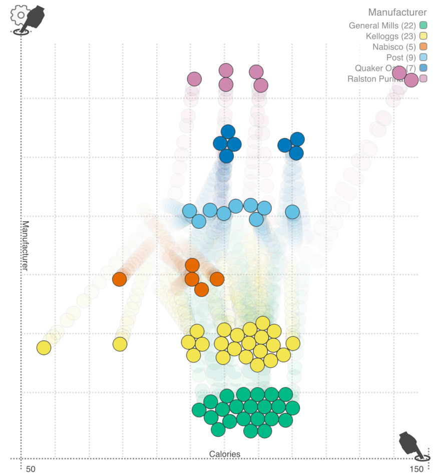

Figure 3: A user can draw a scatterplot using two fingers, one at each corner. When they place the scatterplot, the points feel a pull towards their proper place on the chart. They move towards their location, forming clumps if the points are concentrated in a certain range of values.

As we felt a desire to interact with the data in a way not yet written into the program, we implemented new features. When an affordance went unused or superseded, we removed it. Over many successive passes, we designed two primary types of tools based on the affordances we explicated earlier: manipulative tools and interrogative tools. Manipulative tools alter points' locations, or move them around the sandbox. Interrogative tools change the appearance of a point or its interactivity. A mix of these tools can be layered to explore multiple dimensions at once, and the tools leave traces on the sandbox field so users can see what is affecting points.

In the following sections we describe the physics-based tools implemented in Kinetica through two use cases based on participants' explorations during our user study.

Use Case 1 - Choosing a Car

Margaret would like a new car, but does not know which model to buy. There are many on sale, and each car has a dozen or so different columns of statistics, including weight, fuel economy, and manufacturer. She loads a database of car models into Kinetica, and begins exploring.

She wants a vehicle that is somewhat powerful but uses little fuel. To help answer questions like this, we combined physics-based forces with traditional multivariate charts: histograms and scatterplots. *Histograms* pull points to their proper place on a number line or categorical division. If there are a lot of points in a particular part of the distribution, they will bunch up and visually consume more area. *Scatterplots* functioned similarly, pulling points based on their values in 2D. We initially had concerns of points blocking one another, leading to data being located in an incorrect position on a histogram. To help avoid this overconstraint, we temporarily disable point collision until the data settle down into the proper locations. While this

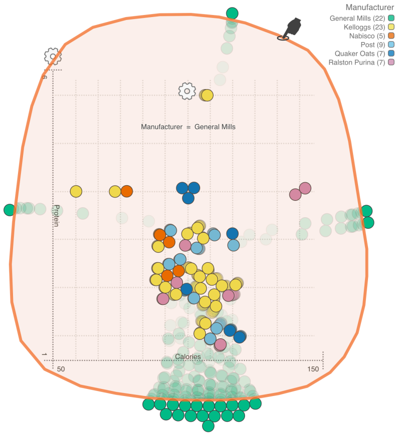

Figure 4: Instead of pulling a barrier across points, a user can also draw a region that filters points. Here a user has chosen to filter points from a region that includes a chart. Once they finish drawing, the points are immediately pushed outwards..

breaks with the physical model, it makes the resulting visualization more accurate and produces a satisfying "settling" effect for the points.

Margaret places a scatterplot that graphs horsepower and mileage, watching as the points are pulled into the chart (Figure 2). Because she sees the points moving, she has an implicit awareness of the action she just made, and can attribute the motion of the data to a specific operation (Figure 3). She sees a big clump of cars with moderate horsepower and mileage as expected, and a longer tail or low horsepower, high mileage cars. Because the points collide, she can interpret their distribution at a glance, and she easily spots the super high efficiency outlier in the lower right corner because it is so visually distinct.

She sees several good candidates in the lower right corner of the chart, but she would like to filter hatchbacks cars. For filtering actions, one could imagine simply removing the points from the screen. However, we wanted to maintain users' mental model of the data landscape, and simply removing points entirely would break it, leading to uncertainty. Instead, we implemented a *semi-permeable barrier* that could be dragged over points or placed onto the field. The points build up against the barrier, showing just how much has been filtered (Figure 1).

We implemented a collision layering system so that points that should pass through do not get mired in the pile of points backing up against the barrier. Each barrier adds a bit to a bitmask, with points permitted to pass the barrier assigned a 1. When points collide, they compare their bitmasks. If they are the not the same, the points can pass through one another. So if there are two barriers, points that can pass through both form one "layer" of collisions, points that can pass through exclusively one barrier form another "layer", and so forth. This creates some edge cases where collision should happen but does not, such as when barriers are nested. In the future we would like to incorporate regions of space into the algorithm to solve this issue.

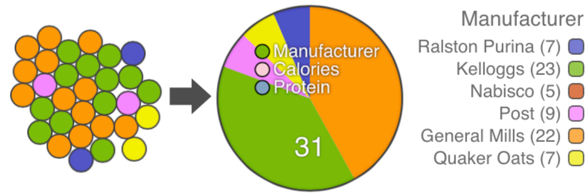

Figure 5: A set of data points can become a group. The pie chart updates dynamically to show distributions.

Margaret drags a wall across the scatterplot, picking up filtered points as she drags. She finally taps to make the wall permanent, separating them off to the right of her screen (Figure 2). As she drags, she notices most of the points she picks up come from the middle of the distribution, leading her to believe most hatchbacks had moderate horsepower and efficiency. The process of filtering not only moves the points away, but also provides insight into relationships between variables and data.

Margaret also has a budget, so she needs to filter expensive cars. However, she doesn't want to move the points since she already know where most of her candidates lay in the sandbox. For this situation, we developed a movable *lens* (Figure 2) that highlights points that match a criterion, performing an analogous, non-manipulative role to the barrier. The lens also can optionally affect the physics within its area, only allowing points that meet its criteria to accept interactions.

Finally, Margaret is ready to investigate individual points. Because all of the tools she used left traces in the sandbox, it is easy for her to backtrack and find the points she remembered were interesting candidates. She uses a double tap to access a *detail view* to explore them. We wanted to avoid occlusion, so the taps place a static magnification glass and show the details of the points within in the glass with labels to help pick out the right one.

Use Case 2 - Titanic Shipwreck

After watching a movie, Margaret is interested in the survivors and victims of the Titanic shipwreck. She is curious if the old adage "women and children first" actually bears out in the data. To test her hypothesis, she loads the data up into Kinetica (see Figure 6). She begins by *coloring* the points based on survival. We implemented separate, color-blind friendly scales for both continuous and nominal values, and also allow users to scale points' size by value.

Margaret now charts the passengers based on gender and cabin class. With a quick glance it is obvious more women than men survived due to the coloring (Figure 6). To investigate how children fared, she draws a semi-permeable barrier around the data, which pushes out points whose age is less than 18. Because the points still feel a pull to their place in the chart, this has the emergent property of keeping the filtered points separated into gender/class groups. From this new chart she makes the conclusion that indeed most children survived, though male children in the third class still faced poor odds.

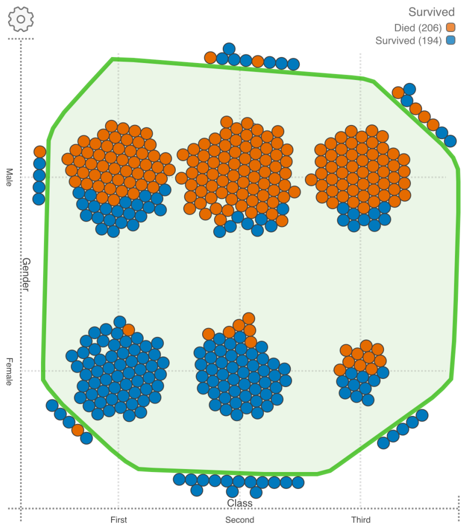

Figure 6: 400 randomly shown passengers are colored by whether they survived (blue) or died (red), and are charted by cabin class (x) and gender (y). Note many more women (top) survived. A filter (green) pushes all passengers under 18 years of age to the outside. Because they still feel a pull to their place on the chart, they separate into class/gender groups. Note third class boys (upper right) still had poor survival odds.

An issue with large datasets is screen real estate. To increase scalability and provide the ability to categorize data, we added a *grouping* function (Figure 5). Users may select a region of points or manually choose points to incorporate into groups. These groups express the characteristics of their constituents, obeying forces as if they are a point whose data is the average of their members.

Gestures and Combinations

For each of the tools we developed, we initially assigned unique gestures and numbers of fingers. For example, to create a barrier, the user would put down four fingers along the contour, while creating a line histogram required only two. This proved cumbersome and confusing. Instead, we adopted gestures that used either one or two fingers. Two fingers define two control points allowing for a histogram between fingers, a bounding box for scatterplots, a barrier between fingers, a lens spanning fingers, or a group that floats between fingers (to avoid occlusion) that selectively consumes points. On the other hand, one finger gives us one control point, which allows for drawing actions. This permits freehand drawn histograms, drawn areas that permit/deny points, drawn lenses, and selecting specific points to form a group.

In testing, we observed several interesting combinations of tools. A scatterplot that bunched points into categorical clusters could be enhanced with a histogram that pulled points into sorted order. Each cluster still felt a pull towards its group, but also self-assembled into sorted order within the cluster thanks to the influence of the histogram. Similarly, drawing an area or barrier that rejected a subset of points while a scatterplot was present meant that even the filtered points still felt a pull to their proper locations, albeit stopped by the barrier. This meant that they still took the form of clusters against the barrier. In practice through techniques like this we could effectively layer 5 dimensions of data within the iPad screen sandbox.

USER STUDY

To evaluate Kinetica, we invited participants to use the software for 45 minutes in a lab study. As a comparison case, we also invited another group of participants to follow the same study protocol using Excel rather than the application. We considered incorporating exploratory visualization suites such as Tableau, Many Eyes, or infogr.am, however we were concerned that training would consume too much time in comparison to Kinetica, especially considering novice computer users. We recruited from a campus participant pool, and scheduled sessions over the course of two days, one for each technique.

We solicited samples for 16 participants in each condition, however one Kinetica participant did not complete the study protocol and their submission was excluded. 15 participants identified as male, and 16 as female. Their average age was 28.9 years. 20 participants reported as using statistical software like Excel at least once a month, and all but one reported having used spreadsheets to analyze data at least once. 11 out of 15 Kinetica users reported having used a tablet, though only 5 owned an iPad. Excel was the only tool participants identified as a way to explore tabular, multi-dimensional data, and all but 2 participants had used it before.

At the start of the study participants received either an iPad with Kinetica loaded or a preloaded instance of Excel with study data present and a populated example chart and pivot table. The study first provided five minutes of basic training. In the case of Kinetica, participants followed a built-in tutorial that goes over each tool with a use case example. For the Excel participants, the study observer presented two well-viewed YouTube video tutorials on Excel and pivot tables/charts. In both conditions study participants were informed that their interface actions were being logged, and an observer watched during the study.

After training, participants were asked to look at a dataset of 73 different brands of cereal containing nutritional information (see Figures 3, 5). They had 10 minutes to use their respective tools to answer a set of 5 basic questions of increasing difficulty/dimensionality. These ranged from “What cereal has the most calories,” to “Of cereals with more than 2g of protein and less than 160mg of sodium which has the most fat?” This also functioned as additional training, and the observer was available to answer questions about how to use the software, but not give hints about the questions themselves. Participants averaged 3.4 questions correct, and there was no difference between conditions. This suggests that perhaps neither condition benefitted more from training, or that the choice of software did not affect participants’ basic statistical ability.

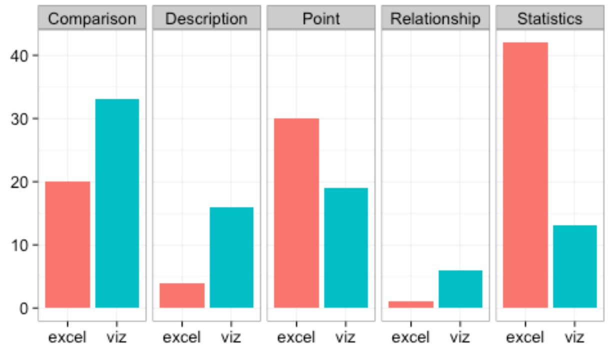

Figure 7: Coded Titanic findings made during user studies

Following the cereal question period, we asked users to pick a car they would like to buy from a dataset of 133 car models with specifications. Purchasing decisions often require people to balance multiple different kinds of data, making them well suited for multivariate visualizations. We first asked users to write down the criteria they would use for buying a car to make them to think about different dimensions of the data. Participants commonly reported considering multiple dimensions at once, including cost, horsepower, gas mileage, and car body type. Afterwards, they used Excel or Kinetica to pore through the data and pick a car they liked. All participants were able to identify at least one car to buy. Kinetica participants used a variety of tools to complete this task. Some plotted points out, then used barriers to filter points. Others layered multiple lenses onto a field, manually pushing points from criterion to criterion. The Excel participants used tools such as charts and pivot tables. Every participant considered at least 2 dimensions of data, with some Kinetica participants considering 4 or 5. This suggests that Kinetica might help participants consider more dimensions of data at once. We will investigate this more quantitatively in the next task.

In the final 15 minutes, we gave participants a dataset containing a random sample of 200 passengers on the Titanic. Since they now had experience with the software, we gave them an open ended prompt. There were to make as many findings as they could in the time limit using the data. For instance, they might identify who the youngest passenger was, or see a relationship between cabin class and passenger survival. Excel participants made on average 5.1 findings, and Kinetica participants made on average 5.5 findings (the difference is not significant). The Kinetica participants' findings generally encoded more dimensions per finding than Excel participants (M:1.74 vs. M:1.38; repeated measures $F(1,164)=17.67$, $p<0.001$).

We were especially interested in whether the nature of participants' findings would differ by condition. In particular, we hypothesized that physics-based affordances would grant participants a more holistic understanding of the data, increasing their awareness of trends and distributions and enabling them to layer and investigate multiple dimensions at once. To investigate this we had two coders classify each of the findings made by participants into five types: *point* findings that discussed a particular row of data (youngest passenger), *statistical* findings that gave summary statistics (mean age), *descriptions* that capture general trends (less than half of the people lived), *comparisons* between categories or groups (more third class passengers died), and *relationships* between dimensions (the older you were, the more likely you were in a high class cabin). Coders were blind to condition and had high interrater reliability (N=186, kappa=0.96).

Interestingly, Kinetica users made far more descriptive, comparative, and relationship findings than their Excel peers whom almost always made point or statistical findings (Figure 7; $c^2(5)=31.3$, $p<0.001$). Thus while they do not make as many quantitative observations, they demonstrate a greater awareness of distribution and multidimensional trends. This suggests that physics-based affordances may be helping users build a more holistic understanding of data.

Qualitative Responses

The qualitative responses during the survey suggest a similar trend. On 7-point Likert scales, Kinetica participants reported that the tool was fairly easy to learn (M:5.53), made it easy to explore data (M:6.0), was fun to use (M:6.07), was fairly easy to use (M:5.07), and that they would use it again (M:6.0). They identified the scatterplots and lenses as especially useful (M:6.2, M:6.47), and groups as less useful (M:4.93), perhaps because they broke with the force-based tools that comprised most of the interface by absorbing points rather than moving/altering them.

After the study, a participant reported that Kinetica was, "… so much better than Excel. I can see when my sort isn't quite right because a point will be where it shouldn't be. It's visual." This points to both the visual memory and process benefits of physics affordances. Another reported, "I like it 'cuz you could just compare things because they were right there. You didn't have to look between rows or anything." In written free response sections, a participant wrote "The scatterplot allows you to visualize your most important priorities [then] …The lens helps you narrow down scope & they can be layered on top of each other," identifying how the layering of physics constraints matched their internal exploration process. Other participants identified ludic elements, stating Kinetica "allows you to "play" around with the data," "makes seeing the data fun and interesting," and "fluidity made it engaging." One participant reacted strongly after the car portion of the study, saying aloud "You can [use Kinetica to]… make really informative decisions on what you purchase in your everyday life."

Participants also had some issues using Kinetica. One participant complained that it was hard to get exact numbers for things. While the interface was well suited for qualitative findings, the participant had to tap three times to see numeric details for a point. Interface limitations did not help, as configuration sliders could not always provide

Forces - Act on points to move, attract, or accelerate
Layouts - Place points into meaningful locations
Mutations - Mutate, combine, or change points
Barriers - Prevent points from occupying a region
Filters - Selectively include or exclude based on criteria
Queries - Change the display appearance of points

	Forces	Layouts	Mutations	Barriers	Filters	Queries & Overlays
Forces	***Magnets*** Points are pulled to a finger					
Layouts	***Force plots*** Points are pulled to their place in a chart	***Fixed charts*** Points snap to their place on a chart				
Mutations	***Group by vector*** Group points that are moving similarly	***Group by layout*** Group points by their place in a chart	***Groups*** Put points into abstract groups			
Barriers	***Jelly walls*** Points need an extra push to pass a wall	***Layout walls*** As points pass, put them into a layout	***Collision to data*** Points that hit a wall are assigned values	***Walls*** Block points from a region of space		
Filters	***Query magnet*** Attract points that meet specific criteria	***Detail view*** See point details in a layout popup	***Group by filter*** Group points based on specific criteria	***Sieves*** Walls only permit certain points	***Show/hide*** Display points by criteria	
Queries & Overlays	***Query fields*** Visualize the pull of different forces	***Area details*** See details for points that meet criteria	***Group zoom*** Zoom in and manipulate points	***Lens*** Highlight certain points in a region	***Selective overlay*** Highlight/select points by criteria	***Overlays*** Color/scale points by data dimensions

Table 1: Generative framework that combines different physics-based affordances to generate new tools/interactions

enough granularity to participants so that they could set exact filtering criteria they wanted for lenses and barriers. Kinetica did not have an undo action implemented, and some participants complained that they wanted to go back, but had already moved data around since then. This raises an interesting question of what undo operations mean in this space. Kinetica also used basic on-screen controls rather than gestures (see buttons in Figure 2). This meant that they sometimes got in the way, annoying users. Of all the tools, participants voiced the magnet as being the least useful, primarily because its one-dimensional force was subsumed by histogram charts whose behavior was clearer.

GENERATIVE FRAMEWORK

Our initial work in Kinetica points to a fertile new area for exploration in data visualization. Physics-based visualization affordances make use of the inherent expertise users have based on their experiences in the everyday world in order to help them develop an understanding of data. These techniques are different from traditional visualization approaches, and, in light of users' desire at times for more familiar controls, may work well in concert. However, because they are different from traditional approaches, it is not always easy or intuitive to create new interactions.

As we explored physics-based visualization techniques in Kinetica, we developed a framework to generate potential affordances. From our development, we propose the following set of physics-based affordance primitives to enhance interactive multivariate data visualization:

- **Data** are represented as physical points that have associated physical properties that correspond to their values in different dimensions
- Data occupy a **sandbox** that contains them and allows for interaction. Interactions with the sandbox change the physical arrangement of the data, and leave traces of their activities.
- The user can employ **forces** to act on physical points either independently or as a result of their unique data. For instance, a magnet may repulse points with low values in a particular dimension.
- The user may use **layout** tools to force points into strict, meaningful locations, breaking with the physics metaphor when necessary (such as when allowing points to pass over or under others to avoid being trapped).
- The user can **mutate** points, for instance combining multiple points into one group so as to observe more points at once or see larger trends.
- The user may place **barriers** that block or selectively block points based on criteria.
- The user can employ **filters** to selectively include or exclude points to help avoid overload or choose only a small subset of interest.
- The user may use **queries and overlays** to change the appearance or behavior of points on the screen.

We imagine these affordances could work in a variety of contexts and situations. They could be used in 2D, which may be easier to interact with (as Sedlmair, Munzner, and Tory [29] suggest in the context of scatterplots), or in 3D where there is a richer space of interactions and there is room for more dimensions. Likewise, they can function using a keyboard, mouse or multi-touch.

Furthermore, these primitives could be combined to generate a much richer set of potential physics-based affordances. Mixing these different techniques together can provoke new ways to augment data visualizations with physics. Table 1 provides some examples of different tools that can come out of a combinatorial brainstorming on these primitives. For instance, combining a barrier that blocks points with a selective filter could generate Kinetica's permeable barrier. Combining forces with a scatterplot layout could create a force-based plot that pulls points to their proper location, and might work in concert with other barriers and forces. We found this framework useful for building the specific physics-based interaction methods we instantiated into Kinetica. We hope it will be useful for future researchers and practitioners in generating even more interaction types that can promote effective physics-based interactions with data.

DISCUSSION AND LIMITATIONS

In this paper we introduced a general set of physics-based affordances for multivariate data visualization. These affordances come with both benefits and costs. For example, on one hand, tools that are easy for users to intuit and interpret might make exploring data easier. On the other hand, constraint-based systems may produce output whose causes are difficult to interpret. In this section we explore these benefits and costs through the lens of Kinetica and more generally, and describe future work that may help to make physics-based affordances more versatile.

Potential Limitations

Scalability. Because of the iPad's limited computational resources and screen real estate, it is hard to analyze more than 500 points without either the rendering being sluggish or the screen filling. While making points smaller and simpler helps, the device itself is a limitation. This hearkens back to the more general scalability limitation of physics-based affordances: some devices, sizes, and interfaces are more suited to certain amounts of data and ways of interacting, and cognitive or attentional resources may be constrained to smaller amounts of data onscreen. While it might only be possible to render 200-500 points on an off-the-shelf iOS, Windows, or Android tablet, those 200 points could be abstract representations of thousands of points. Users might zoom in and out of different levels of detail (or infinite levels of detail [4]), especially given systems that help users make approximate queries and abstractions [8]. As devices grow in power and can render more points, this need may diminish, though the constraints of human visual attention and sensemaking might still suggest limiting the number of points shown.

Another possibility is to use different devices for different purposes. While a large surface may display all points, a smaller device in tandem might help users explore a smaller subset using more detail-oriented interactions. The large- and small-scale tools may not even need to share the same tool set because of their different needs. Instead of using differently sized displays, one might also employ collaboration. One user might take one section of a larger continuous physics environment, handling a certain kind of points, while another manages a different area.

Effect attribution. Participants in the study also expressed several situations where they had a hard time attributing effects. If they placed too many force-based plots on the screen, it became a mess of points that was hard to interpret. This is a general problem inherent in physics-based affordances. As more constraints and forces interact in the sandbox, it becomes harder to attribute them to specific causes. In the case of Kinetica, participants quickly found a solution: using the visual traces of what forces are currently in use to backtrack until the sandbox made sense again. It is possible that more feedback could help solve this issue, but could necessitate a tradeoff with increased clutter. For example, in a case where three forces are interacting on a small set of points all at once, perhaps each force could produce waves or lines of force that indicate their effects on the point.

Expert usability and overload. A few participants voiced concerns similar to Drucker et al. [6], where their expertise and familiarity with traditional WIMP statistical tools made them want quantitative tools like means or the ability to draw traditional bar charts. This is also suggested by Excel participants' ability to make more point-specific or quantitative findings compared to Kinetica. In principle such statistics could be added to a physics-based visualization, for example in a panel, though again tradeoffs in terms of training and screen real estate may be incurred.

Overconstraining. One risk in any constraint-based system is that points become trapped by the constraints such that the visualization does not faithfully represent the underlying relationships in the data. We tried very hard to eliminate such situations during the design of Kinetica. We introduced measures that broke with physical consistency as needed, allowing points to pass through one another to avoid bunching up at permeable barriers. If points are being pulled into a sorted order, then we once again disable collision temporarily to prevent mis-sorts. While such breaks with physical models could reduce the consistency of the visualization interactions, they represent a tradeoff between usability and data faithfulness. This is a general problem that may affect physics-based affordances as they are developed further, exacerbated by the difficulty of predicting what minima or overconstrained situations might occur in practice.

Potential Benefits

Minimal training. Participants approached our novel interface, some never having used a touch device or tablet before, and were able within ten minutes to explore and generate findings from data they had never seen before. As Jacob et al. [18] mention, users already have an intuitive awareness of physics and their surrounding environment. As a result, many of these primitive affordances are in a sense already "trained" [27]. Users know that a magnet pulls objects towards it even if they do not know about magnetic forces and inverse square decay. This minimal barrier to entry surprised participants. One suggested that these approaches may be useful for teaching data literacy to young students or visual learners.

Outlier detection. Participants regularly spotted outliers in the data because they were in unusual locations, had different colors than the points near them, or moved unlike the rest of their nearby points. These properties emerged out of their general use of physics-based tools, and spatial correlations and groupings often make identifying outliers easier [14]. Participants felt encouraged to explore these points and try and explain why they were different, generating stories and expressing surprise.

Awareness of distribution. Kinetica participants made more descriptive and comparative findings, describing amounts of points and using words like *more* or *less*. Because data in a physics-based visualization are treated as physical objects, they can bunch up and form salient visual patterns. Imagine a histogram or scatterplot where points collide with each other (see Figure 2). Places where there is higher density of data are inherently more bunched together. This fits users' intuitions, and gives an implicit, visual depiction of data distribution without a need to chart frequencies. However, this qualitative understanding may come at the expense of the quantitative, point-by-point understanding that Excel can provide. A hybrid of these approaches, with more quantitative analytical techniques stepping in when a user is ready to focus on a small subset of data might give users a satisfying middle ground.

Understanding process. Kinetica participants also cited a better understanding of process in their qualitative responses. They mentioned fluidity and how everything was "right there" when they needed it on the screen. When participants used Kinetica to explore their data, the physics interactions unfolded in front of them, with points moving to their proper places on charts and gathering behind permeable barriers. These small changes, over time, helped them to interpret the end state of an operation and gave them new ideas for future hypotheses to explore. This may be a more general property of physics-based affordances, since as groups of points move together and possibly share visual features, people can gain a summary/qualitative understanding of their data [14].

Augmenting working memory. Participants in the Kinetica case also generated findings that made use of more dimensions than their Excel peers. This may be because of an ability to consider and keep track of more dimensions. As a user combines interactions in the process of exploring multivariate data, they leave traces of their explorations encoded onto the sandbox. For example, if a user uses a barrier to separate points based on one dimension, then the barrier remains as visual element, and the points are physically separated into two different areas of the sandbox. Because each of the tools leaves traces in the sandbox (see Figures 2, 3), it is easy to trace back what sorts of dimensions are at play. Further, because the user has been slowly moving points and arranging the field, the sandbox encodes contextual cues for the user which can improve their memory of the current state of the data and their past actions [5].

Future Work

Kinetica only took a small slice of the physics-based affordances we generated using our framework. It could make use of many more, perhaps solving some user concerns in the process. In its current state, Kinetica does not do any layout computations. Everything is force directed, which runs the risk of local minima or overconstrained blocking. Layout algorithms that put points into predefined, semantically meaningful configurations might help users make sense of new data and also help solve some overload. For instance, points could be forced into clusters decided by k-means clustering, or pushed into tabular rows that are easy to manipulate along side a richer query language. Other than just filtering points by a simple < / = / > criterion, barriers might employ complex logic or themselves push points that pass through into layouts.

More generally, physics-based affordances are not an end-all solution for multivariate data visualization. In Kinetica we mixed in traditional scatterplots because they were so familiar and so useful. Such affordances probably work best in a hybrid situation. If there are millions of points, maybe a WIMP technique on a computer is best for a first pass. Even within a system like Kinetica there also may be a need for traditional visualizations like parallel coordinates and plots. While this may break the consistency of the data-to-point physics model, it may improve user capabilities, especially if they already have expertise.

CONCLUSION

In this paper we contributed an outline of physics-based affordances to augment multivariate data visualization. Using this outline, we developed a generative framework for ideating new interactions and affordances in this space. We instantiated some of these different affordances in the form of an iPad application, and conducted a comparative study between the physics-based visualization tool and a traditional WIMP spreadsheet tool with charts and pivot tables. In general the physics-based tool was more satisfying for participants, and led them to a more holistic understanding of the data they were exploring. Although this work is only one step towards a future of enhanced data interaction, we hope that in the future an interdisciplinary group of interface, information visualization, human factors, and other researchers will generate new and exciting approaches for exploring data that harness the growing tide of new devices and new ways of interacting.

ACKNOWLEDGMENTS

This work was supported by NSF grants IIS-1149797, IIS-1217559, IIS-0968484, IIS-1111124, and Microsoft.

REFERENCES

1. Agarawala, A. and Balakrishnan, R. Keepin'it real: pushing the desktop metaphor with physics, piles and the pen. *CHI*, (2006).

2. Ahlberg, C. and Shneiderman, B. Visual information seeking: Tight coupling of dynamic query filters with starfield displays. *CHI*, (1994), 313.
3. Baur, D., Lee, B., and Carpendale, S. TouchWave: kinetic multi-touch manipulation for hierarchical stacked graphs. *ITS*, (2012).
4. Bederson, B. and Hollan, J. Pad++ : A Zooming Graphical Interface for Exploring Alternate Interface Physics. *UIST*, (1994).
5. Chun, M. and Jiang, Y. Contextual cueing: Implicit learning and memory of visual context guides spatial attention. *Cognitive psychology 36*, 1 (1998), 28–71.
6. Drucker, S., Fisher, D., Sadana, R., and Herron, J. TouchViz: a case study comparing two interfaces for data analytics on tablets. *CHI*, (2013).
7. Elmqvist, N., Stasko, J., and Tsigas, P. DataMeadow: a visual canvas for analysis of large-scale multivariate data. *Information Visualization 7*, 1 (2008), 18–33.
8. Fisher, D., Popov, I., Drucker, S., and Schraefel, M. Trust me, i'm partially right: incremental visualization lets analysts explore large datasets faster. *CHI*, (2012).
9. Frisch, M., Kleinau, S., Langner, R., and Dachselt, R. Grids & guides: multi-touch layout and alignment tools. *CHI*, (2011).
10. Frisch, M., Langner, R., and Dachselt, R. Neat: a set of flexible tools and gestures for layout tasks on interactive displays. *ITS*, (2011).
11. Fruchterman, T., and Reingold. E. "Graph drawing by force‐directed placement." *Software: Practice and experience* 21, 11 (1991), 1129-1164.
12. Goth, G. Brave NUI world. *Communications of the ACM*, (2011).
13. Hancock, M., Cate, T., and Carpendale, S.. "Sticky tools: full 6DOF force-based interaction for multi-touch tables." *ITS*, (2009).
14. Haroz, S. and Whitney, D. How Capacity Limits of Attention Influence Information Visualization Effectiveness. *IEEE Transactions on Visualization and Computer Graphics 18*, 12 (2012), 2402–2410.
15. Huron, S., Vuillemot, R., and Fekete J. "Visual Sedimentation." *IEEE Transactions on Visualization and Computer Graphics*19.12 (2013): 2446-2455.
16. Inselberg, A. and Dimsdale, B. Parallel coordinates: a tool for visualizing multi-dimensional geometry. *IEEE Transactions on Visualization and Computer Graphics*, (1990), 361–378.
17. Isenberg, P. and Isenberg, T. Data Visualization on Interactive Surfaces : A Research Agenda. *IEEE Computer Graphics and Applications*, April (2013).
18. Jacob, R., Girouard, A., Hirshfield, L., et al. Reality-based interaction: a framework for post-WIMP interfaces. *CHI*, (2008).
19. Jacob, R.J.K., Ishii, H., Pangaro, G., and Patten, J. A tangible interface for organizing information using a grid. *CHI*, 1 (2002), 339.
20. Keefe, D. and Isenberg, T. Reimagining the Scientific Visualization Interaction Paradigm. *Computer*, May (2013), 51–57.
21. Klemmer, S., Hartmann, B., and Takayama, L. How bodies matter: five themes for interaction design. *DIS*, (2006).
22. Kosara, R., Bendix, F., and Hauser, H. Parallel sets: interactive exploration and visual analysis of categorical data. *IEEE Transactions on Visualization and Computer Graphics 12*, 4 (2006), 558–68.
23. Lee, B., Isenberg, P., Riche, N.H., and Carpendale, S. Beyond Mouse and Keyboard: Expanding Design Considerations for Information Visualization Interactions. *IEEE Transactions on Visualization and Computer Graphics 18*, 12 (2012), 2689–2698.
24. Lee, B., Kazi, R., and Smith, G. SketchStory: Telling More Engaging Stories with Data through Freeform Sketching. *IEEE Transactions on Visualization and Computer Graphics*, August (2013).
25. Lembcke, S. Chipmunk Physics. www.chipmunk-physics.net.
26. Martin, A. and Ward, M. High Dimensional Brushing for Interactive Exploration of Multivariate Data. *VIZ*, (1995).
27. North, C., Dwyer, T., Lee, B., and Fisher, D. Understanding multi-touch manipulation for surface computing. *INTERACT*, (2009), 236–249.
28. Schmidt, S., Nacenta, M., Dachselt, R., and Carpendale, S. A set of multi-touch graph interaction techniques. *ITS*, (2010), 113.
29. Sedlmair, M., Munzner, T., and Tory, M. Empirical Guidance on Scatterplot and Dimension Reduction Technique Choices. *IEEE Transactions on Visualization and Computer Graphics*, August (2013).
30. Shneiderman, B. Dynamic queries for visual information seeking. *IEEE Software*, (1994), 1–18.
31. Spitzer, K., Eisenberg, M., and Lowe, C. Information Literacy: Essential Skills for the Information Age. *DESIDOC Journal of Library & Information Technology 28*, 2 (1998), 39–47.
32. Stasko, J., Goerg, C., and Liu, Z. Jigsaw: Supporting Investigative Analysis through Interactive Visualization. *Information Visualization 7*, 2 (2008), 118–132.
33. Ullmer, B., Ishii, H., and Jacob, R. Tangible query interfaces: Physically constrained tokens for manipulating database queries. *INTERACT*, (2003).
34. Walny, J., Lee, B., Johns, P., Henry Riche, N., and Carpendale, S. Understanding Pen and Touch Interaction for Data Exploration on Interactive Whiteboards. *IEEE Transactions on Visualization and Computer Graphics 18*, 12 (2012), 2779–2788.
35. Wilson, A., Izadi, S., Hilliges, O., Garcia-Mendoza, A., and Kirk, D. Bringing Physics to the Surface. *UIST*, (2008).
36. Wobbrock, J.O., Morris, M.R., and Wilson, A.D. User-defined gestures for surface computing. *CHI*, (2009), 1083.
37. Wong, P. and Bergeron, R. 30 Years of Multidimensional Multivariate Visualization. *Scientific Visualization : Overviews, Methodologies, and Techniques*, (1994), 3–33.
38. Wu, M. and Balakrishnan, R. Multi-finger and whole hand gestural interaction techniques for multi-user tabletop displays. *UIST*, (2003), 193–202.
39. Yi, J.S., Melton, R., Stasko, J., and Jacko, J. a. Dust & Magnet: multivariate information visualization using a magnet metaphor. *Information Visualization 00*, April (2005), 1–18.
40. Zickuhr, K. *Pew Report: Tablet Ownership 2013*.