Multi-Step Animation to Facilitate the Understanding of Spatial Groupings: the Case of List Comparisons

Catherine Plaisant[*], Tiffany Chao[*#], Ran Liu[*#], Kent Norman[*&], Ben Shneiderman[*#]
[*]Human-Computer Interaction Lab
[#]Dept. of Computer Science, [&]Dept. of Psychology
University of Maryland, College Park, MD
{plaisant, ben}@cs.umd.edu

ABSTRACT

While animation has been shown to be compelling and helpful to reveal transformations of complex graphical representations such as trees or graphs, other studies have cast doubts on animation's usefulness for learning. We present a new beneficial use of animation: helping users learn and understand the meaning of the spatial grouping of items on the screen. We introduce this technique in the design of two list comparison interfaces: Twinlist, an interface that helps physicians compare and merge two separate lists of medications into a reconciled list; and ManyLists, an interface for product comparison. Animation is used to reveal the similarities and differences between items in the lists and explain the final grouping. A controlled experiment confirmed that animation helped participants learn the groupings of Twinlist. Finally we summarize design guidelines and discuss other possible uses of the technique.

Author Keywords

animation, visualization, comparison

ACM Classification Keywords

H.5.2 Information Interfaces and Presentation: User Interfaces—Graphical user interfaces (GUI)

General Terms

Human Factors; Design; Measurement.

INTRODUCTION

Perceiving and interpreting motion is a fundamental element of human perception. Our eyes are attracted to moving objects in the real world and items moving in the same direction are interpreted as a group [Ware, 2012]. Unsurprisingly animations have been a staple of interactive applications since displays have been fast enough to allow smooth transitions. They serve a variety of communication purposes [Novick 2011, Tversky 2002; Mayer, 2009]. Done well, they are extremely compelling, direct attention, and make complex transitions intelligible. Done poorly, they distract, irritate, and waste time. Multi-step animation can be helpful to reveal transformations of complex graphical representations such as trees or graphs, but more examples of successful applications are needed to inspire designers, and additional user studies are needed to further our understanding of which situations benefit from animation and which ones do not.

HCIL Tech Report (2012)
http://www.cs.umd.edu/hcil/pubs/tech-reports.shtml

To this effect, we present a new beneficial use of animation. We believe that multi-step animation is useful to help users learn and understand the meaning of spatial groupings representing similarities and differences between lists of items. We describe two practical uses of the technique:

1) Twinlist (Figure 1), an interface that helps physicians compare and merge two separate lists of medications into a reconciled list. The animation reveals the similarities and differences of the drugs in the two lists.

2) ManyLists, an interface to compare similar products and reveal the identical and unique features of the products, then differentiate the similar features.

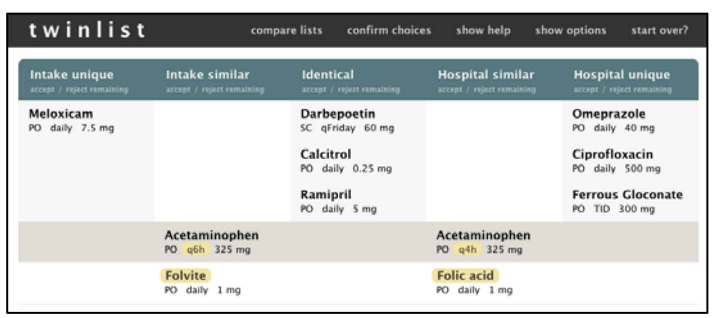

Figure 1. Twinlist uses spatial groupings to convey the similarities between two lists of medications. A multi-step animation was used to help users understand the groupings.

We summarize the results of a controlled experiment that confirms that animation helped participants learn the groupings in Twinlist, and report on a usability study of ManyLists. Finally, we propose design guidelines and discuss other possible uses of the multi-step animation to explain spatial groupings.

RELATED WORK

Many have tried to empirically measure the effect of animation, for example [Bederson et al. 1999] found that animation helped users build mental maps of spatial information. Animation has been shown to be effective in revealing transformations of complex graphical representations such as trees or graphs [Robertson et al. 1991, 2002] [Shanmugasundaram et al. 2007] [Yee et al, 2001], or revealing trends over time [Robertson et al. 2008] [Rosling, 2006]. [Heer and Robertson 2007] report the importance of animated transitions for revealing changes between statistical data graphics. The pacing of the animation has been less studied: [Dragicevic et al 2011] studied object tracking performance under different conditions of temporal distortion during animation, and slow-in/slowout outperformed other techniques, but there were differences depending on the type of visual transition. Still, animation is hard to design and not always helpful. [Tversky et al. 2002] and [Mayer 2009] report on various studies that cast doubts on animation's usefulness. Done well, it is compelling, directs attention, and making even complex transitions intelligible e.g. [Chevalier et al. 2010]. Done poorly, it can also distract and irritate; and waste time. [Fisher 2010] summarizes both viewpoints. [Lawrence et al. 1994] showed that animation was helpful to students learning algorithms when they built the animation themselves. [Tversky et al. 2007] summarized with three animation aphorisms: "Seeing isn't perceiving; perceiving isn't understanding; showing isn't explaining".

Animations might include one or multiple steps (also called transitions, sequences, frames, or stages). Long, canned animations are similar to videos and typically include many steps. They have been used to explain complex processes such as manufacturing or to help users get started with complex interfaces using recorded demonstrations [Plaisant 2005]. On the other hand, most early animations used in interactive systems involved a homogeneous, continuous animation (e.g. fisheye effect, continuous 3D interaction) or cartoon inspired representations [Thomas 2001] that did not necessarily match the actual operations (the trashcan opened and closed to swallow a file icon). Small animations now pepper elegantly designed applications (e.g. many mobile apps) to create an attractive look and feel.

An early example of step-by-step animation aimed at explaining the placement of items on the screen was Spacetree [Grosjean et al 2002] which decomposed the animation of a changing tree representation in three main steps: trim, translate, and grow. Later [Herr 2007] also recommended the use of staging for complex transformation of statistical graphics (e.g. separating axis rescaling from value changes). [Guilmaine 2012] propose to give users control over hierarchically organized animation steps. Backing up those designs [Tversky 2011] suggests that step-by-step animation is more congruent to the way the mind understands and represents continuous organized action than continuous animation. After reporting on usability studies showing that animation can be confusing in showing trends [Robertson et al. 2008] recommend that presenters ensure the data tells a "clean story" and avoid having too many data points moving at once, data points that reverse their tracks over time, or data points that do not move in synchrony.

In summary, while multi-step animation has been shown to be possibly helpful to reveal transformations of complex graphical representations, new examples and further user studies are needed to sharpen our understanding of which situations benefit from animation and which ones do not.

Next, we describe two list comparison applications that illustrate a new use for multi-step animation: helping users understand the meaning of spatial groupings.

EXAMPLE APPLICATION #1: TWINLIST FOR MEDICATION RECONCILIATION

We used multi-step animation in the design of Twinlist, an interface that helps physicians during the important and complex task of medication reconciliation [Poon 2006 and 2009, Pronovost 2003, Claudino, 2011].

Let's focus on a single scenario of use: discharging a patient from the hospital (Figure 2). Imagine a patient who goes to the hospital for an emergency. When he is admitted to the hospital a nurse compiles the list of medications he was taking at home (through an interview or access to pharmacy records or other electronic health records). This is called the "intake list". Later, during the hospital stay, another list is created to keep track of the medications the patient receives while at the hospital: this is the "hospital list". When the patient is discharged from the hospital, a physician has to compare those two lists and prepare a reconciled list with all the drugs the patient should be taking after his return home.

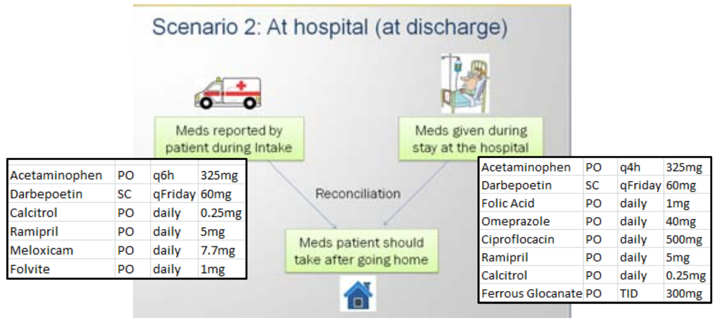

Figure 2. A scenario of medication reconciliation: discharging a patient from the hospital requires physicians to compare two lists of medications: the "intake list "and the "hospital list" . What is identical? unique? similar?

In current interfaces, physicians might see one list in one tab window, the other list on a second window and they have to re-enter the final list in yet another window. Some systems present a single merged list listing all drugs which at least brings close together the drugs with the same name. [Markovitch 2011] reviews the research, describes different levels of drug equivalence and shows that revealing equivalences drugs can simplify the reconciliation task.

Preprocessing

The preprocessing phase in Twinlist identifies equivalent drugs found in the two lists using form equivalence (e.g. Tynelol is just another name for acetaminophen or paracetamol) and functional equivalence (Atenolol and propranolol are both betablockers) [Bozzo Silva et al., 2011]. We then use three categories: drugs are considered "identical" when the same drug appears on both lists (with matching name, dosage, route and form), "unique" when they appear in one list only, and "similar" when the drug has some level of equivalence with another drug in the other list, e.g. same name but different dosages, or same dosage but acetaminophen instead of Tynelol.

Spatial groupings

Twinlist then places drugs on the screen using a multi-column spatial layout (Figure 1). We believe that the spatial groupings help Twinlist provide an intuitive way for users to quickly differentiate items that are the same from those that differ between the two lists. The left half of the screen is for the drugs of the intake list, and the right half is for the drugs taken at the hospital. In the center column we place the identical drugs (i.e. present in both lists: Darpeboetin, Calcitrol and Ramipril). On the far left are listed the drugs unique to intake (here only Meloxicam), on the far right the drugs unique to the hospital. Below this set of three lists we place the drugs that are similar, aligned to facilitate comparison. For example acetaminophen is present in both lists but the frequency of use is different (q6h instead of q4h) so both medications and their details are aligned in the same row, with the difference highlighted in yellow. Folvite is a brand name for folic acid so both drugs are also aligned on a common row so the physician can pick one. The physician has to decide which of the similar drugs is best to use. Showing the origin of the drug is important (intake or hospital).

Armed with the similarity information, physicians can rapidly make decisions to keep or stop drugs, one at a time or entire columns at a time. When they are done, they click "Review and Sign Off" to submit their reconciled list.

Multi-step animation

We use the multi-step animation to help users understand the groupings of drugs (Figure 3).

When the lists are loaded in Twinlist, they are listed side by side: intake on the center left, hospital on the center right. The animation steps are as follows (Figure 3):

1) identical drugs move to the center column, in-between the original lists, and then merge, one pair at a time;
2) unique drugs move away from the center to their respective side, first to the left for the drugs unique to the intake list, then
3) to the right for the drugs unique to the hospital;
4) similar drugs are aligned and golden-yellow highlights are added to indicate the differences between similar drugs;
5) the display is compacted to save vertical space by stacking identical and unique drugs at the top of their respective columns and sliding the rows of identical drugs together below. Options are available to change the speed of the animation or turn it off.

An early prototype written in Adobe Flex [Claudino, 2011] was tested with 2 physicians and 2 graduate students, providing supportive feedback and suggested improvements, which led to a complete rewrite of the software using JavaScript and HTML5, as shown in the figures.

EVALUATION #1

Investigating the value of animation in Twinlist

In order to evaluate the role of multi-step animation in Twinlist we conducted a within-subjects experiment with 20 students (6 male and 14 female, aged 19 to 26.) None had prior experience with either Twinlist or medication reconciliation. All participants used two versions of Twinlist (order was counterbalanced). The "multi-step" version included the full five-step animation while the other version jumped directly to the final layout.

Participants were first given the necessary general background about drug reconciliation and provided with two simplified rules to reconcile the drugs: 1) always select the lower dosage among similar items, 2) reject herbal supplements, probiotics and vitamins (which were labeled clearly as such). Those rules required them to read and compare all the drugs to reconcile the lists. We provided no interface training. We asked them to think aloud while they were exploring and learning the interface on their own with a training dataset, and to tell us when they thought when they were done learning the interface. They were encouraged to say what they found confusing, and ask clarification questions. We recorded the time to learn and the questions they asked. We then provided complementary training when needed, and preceded to the performance task. Participants were asked to reconcile four pairs of drug lists and then asked to provide subjective ratings about the interface. After a rapid training with the second interface they repeated the performance task with that interface. Eight timed datasets varied in appearance and length, and were presented in random order. All participants were tested with a 13.3" laptop, a corded three-button mouse, and Google Chrome.

Learning. We hypothesized that participants would learn to use the interface faster using the multi-step version. A

Welch two-sample t-test found no significant difference in training times ($p = 0.42$). On the other hand we observed differences in terms of comments and clarification questions. For example only 3 of the 10 participants who learned with the multi-step animation reported being initially confused about the five-column layout (compared to 9 of 10 who learned with the one-step version). Fourteen out of 20 stated that they favored learning with the full animation, citing its ability to "show you where everything goes" and what everything is "connected" to. A paired t-test for the related questionnaire item indeed indicated that the full animation was considered more helpful for learning ($p = 0.02$).

Performance. We had hypothesized that participants would make sense of the data while the animation was playing and therefore would start taking action earlier, hopefully reducing the "time penalty" added by the animation. This was not verified: Paired t-tests found that tasks using the multi-step animation took longer ($p = 0.05$) - with the additional time roughly equal to the time of the animation, and that the times to first click and error rates were not significantly different ($p = 0.26$ and $p = 0.22$, respectively).

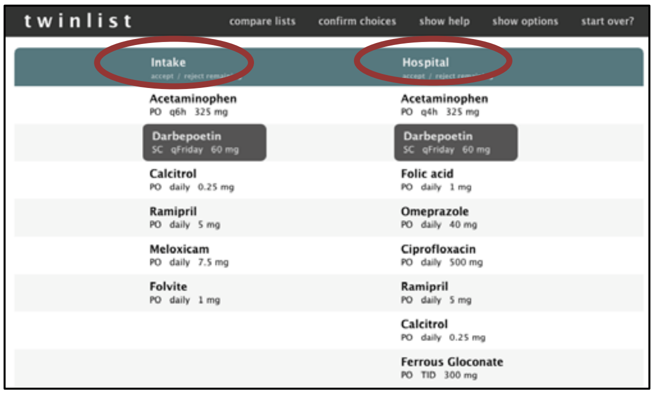

Start: Original layout: two separate lists

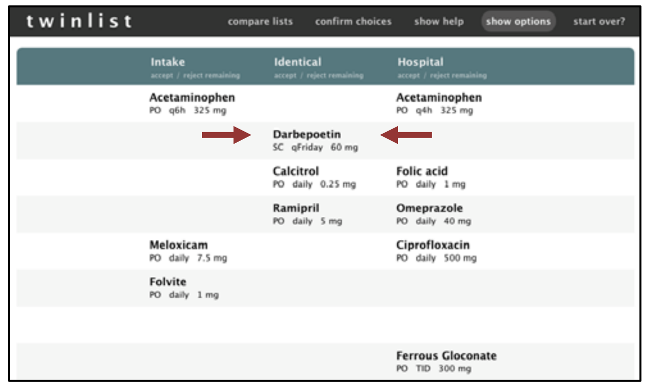

Step 1: The identical drugs move to the middle , one at a time.

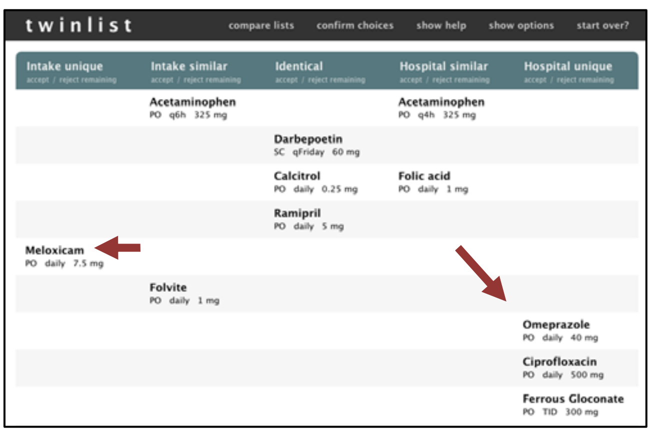

Step 2 and 3 : The Unique drugs move to the left, then the right

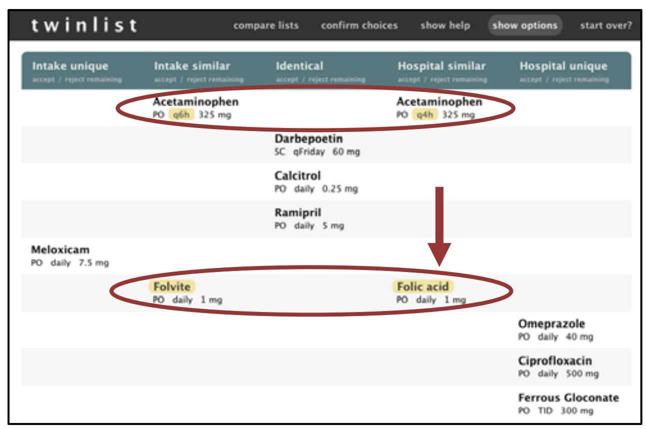

Step 4: The similar drugs are aligned, with their differences highlighted (here dosage differences)

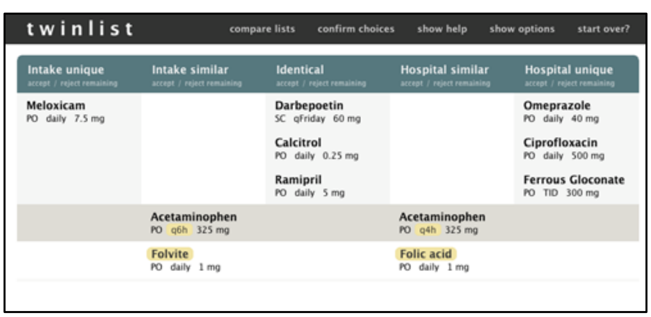

Step 5: The entire display is compacted

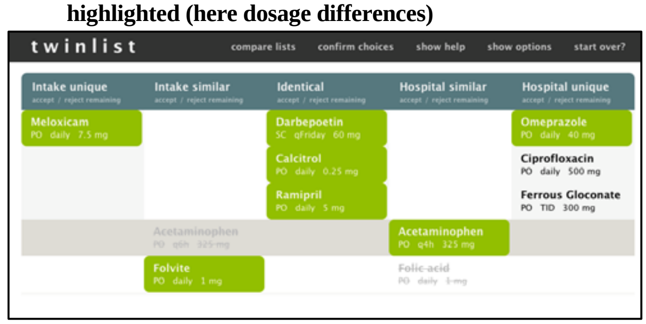

Users select drugs to be kept (green) or not (grayed and stroked out)

Figure 3: The 5 steps of the animation sequence used to explain the spatial groupings (SEE VIDEO FIGURE)

Nevertheless, we observed 7 out of 20 participants clicking on drugs to continue or discontinue them before the multi-step animation had ended, suggesting that indeed they could process information while the multi-step animation was playing.

Preference: We anticipated that participants would prefer the multi-step animation to learn the interface but would prefer to stop the animation for regular use. This was confirmed: As noted earlier Fourteen out of 20 stated that they favored learning with the full animation. Eighteen out of twenty participants stated they would prefer the single step version for regular use (i.e. after learning); a paired t-test of the related satisfaction questionnaire prompt supported this preference ($p < 0.01$).

Lessons learned: Comments collected from participants illustrate the main finding: the multi-step animation "helps you see" initially, but is "too slow" once you know how things work. Some disagreed: two strongly preferred the multi-step animation both for learning and prolonged use, and four strongly preferred the one-step version for both cases. Leaving it as an option that users can turn off is a reasonable design, and since some may not even watch the animation, the column titles should be as clear as possible on their own.

Unsurprisingly, scrolling can cause problems. Five out of twenty participants forgot to scroll the longer lists before attempting to sign off on their choices; another was a near miss. Fortunately, Twinlist handles this type of problem by disallowing signing off until all items have been acted upon.

Other exploration behaviors were observed: for example nine users attempted to drag items, eight out of curiosity, and one to create a "master list". What dragging would be suited for, however, is unclear in this context, but this behavior may have been triggered by watching the animation move items on the screen.

Limitations: The participants in the study were students and not physicians. This meant that learning also including the time to understand the rules of reconciliation. Even though the rules were simple participants did ask clarification questions about those rules and this may have impaired our ability to tease out differences in the time it took to learn the groupings with and without animation. Also, it was easy to see that drugs were similar in our tasks (i.e. the names were quasi-identical, and mostly dosages varied so that participants didn't have to know about real drugs) so users could easily see that aligned drugs were similar. This is true for physicians as well in general as they know many drugs, but in the rare case where physicians would use Twinlist for the 1st time and the only similar drugs aligned were not known to them they may not recognize them as similar. We have not seen this happen in our many presentations to physicians. While we do not have results from a formal study, the reception for Twinlist among clinicians has been enthusiastic. Comments suggest that the animation is helpful, and that the groupings are meaningful. This led to a quick implementation in Microsoft Amalga, an adaptation for a similar problem at Massachusetts General Hospital, and two ongoing projects to add Twinlist to existing EHR systems. A study measuring clinical benefits of the entire interface is underway.

EXAMPLE APPLICATION #2: MANYLISTS FOR PRODUCT COMPARISON

While Twinlist handled two lists of items with a limited numbers of attributes (dosage, route and frequency), ManyLists deals with more than two lists and many more attributes. Examples of comparison situations related to consumer products guided our design:

- Vitamin supplements: here the features are lists of ingredients, available in a standard formats and measures. In almost every case having more of an ingredient is better.
- Food products: they have nutritional ingredients in mostly standard formats. For some ingredients more is better (e.g. fiber, potassium) but for other ingredients less is better (e.g. cholesterol, sodium). For some ingredients user preferences may vary (e.g. while calories are usually considered bad in wealthy countries, sometime consumers want more calories, or more sugar)
- Technical products: for example laptops have many features to compare, but formats are not standardized. Sometimes different brands might have different format for specifications. Numerical values such as hard disk size or RAM are comparable, but other features described might be harder to compare.

Each feature of the products can be classified as being either the same for all products, unique to a single product, or similar across multiple products. Because of the larger number of lists and attributes, a different layout is necessary. Features that are identical (the same) in all products are grouped together and moved on the top of the screen, listed in an horizontal fashion to save space. Features that are unique to a product are listed at the bottom of the screen. Features that exist in two or more products are aligned to facilitate comparison, and grouped in the center of the screen. Highlighting is used to compare the products along those similar features. Here we use a light green color to show the best product among those that have the same feature (e.g. the lowest sugar content).

While we believe this layout can facilitate comparison, it is not a conventional layout of product features. In order to help users understand this layout, ManyLists uses step-by-step animation to help users learn and understand the layout. ManyLists starts off with a standard multicolumn table view where all the features of a product are put in the same column, e.g. four columns for the four different types of coffee drinks in Figure 4. The nutritional ingredients are listed below each product. Five animation steps guide users from the initial layout to the final layout (see Figure 4):

1. The identical features fade out from their initial location as they move to their location in the top "identical" section, one feature at a time (i.e. vitamin A, then caffeine)
2. Unique features drop down to the bottom of their columns in the lower "unique" section
3. The similar elements are aligned, remaining in the middle section of the screen
4. The empty rows are reclaimed to compact and save screen space
5. The legend for the color-coding appears then all the best values are highlighted in green, while the others are grayed out slightly.

Users can use the control panel to replay the animation all at once or step by step. They can also alter the way the features are grouped. For example, since a small change in nutritional value can be considered as insignificant, they can increase the threshold of similarity used to classify features as identical or not, therefore reducing the number of features to compare, e.g. they can specify that products varying in calories count by less than 10% should be considered as identical. ManyLists also supports users changing the default settings of goodness of the feature, e.g. users who seek to lose weight will favor products with fewer calories, while users who are trying to gain weight will prefer more calories.

ManyLists uses a different layout of spatial grouping than Twinlist to reveal similarities and differences, in order to cope with larger numbers of lists and items. Yet a similar multi-step animation can be used to explain the groupings.

EVALUATION #2

To observe how users learn ManyLists we conducted a usability study with fourteen participants (six female and eight male from the university). Thirteen were graduate students. None of the participants were familiar with the ManyLists interface. We collected feedback on how the animation helped participants understand the novel spatial layout, and also observed what strategies users chose to compare the products. We used the coffee dataset shown in the Figure 4 because it is a common type of task, yet moderately complex.

Procedure

We did not provide any training. We described the general context (selecting a coffee drink) and encouraged participants to think aloud while performing the tasks. We then observed the participants' behavior and recorded their comments during and after the tasks.

After they watched the animation, and described their initial understanding of the interface, we asked them to point to the unique features, identical features and similar features among the four products. The next task was to find their preferred coffee out of the four. We observed how users compared products, and if our visual design was found satisfying for them in the decision process.

Users were debriefed with a series of questions regarding their overall impression of ManyLists, the features they liked/disliked, and whether they would choose to use ManyLists when they need to buy a product. We also asked users to make suggestions and rate the design for ManyLists.

Results

Eleven of fourteen of our participants were able to understand the spatial groupings after playing the animation only once. Two of the three who were confused by the animation at first understood it after replaying it a second time. One participants needed us to replay the animation step-by-step, and needed an explanation of each step. Five of fourteen participants complained that the animation for identical features was too fast. One of them suggested we could blink or highlight all the identical features first, and then match them together, so they could have more time to respond to what was happening. Three of them were confused by the highlighted features at first, but after they carefully examined the first features and their values, they all understood the meaning of the green highlighting.

During the debriefing session, typical comments included: "very cool", "I like the animation", "it is really useful", "it is good that it automatically takes out the identical and unique features", "it is very clear", and "I like that you suggest the best values for different features". Common suggestions regarding to the animation included: 1) slow down the animation speed.

Finally, when we asked users to rate the overall impression of ManyLists on a 1-9 scale, where 1 is failing and 9 is exemplary; they gave an average rating of 7.7.

Initially each coffee drink is in a separate columns (i.e. a way it is traditionally presented).

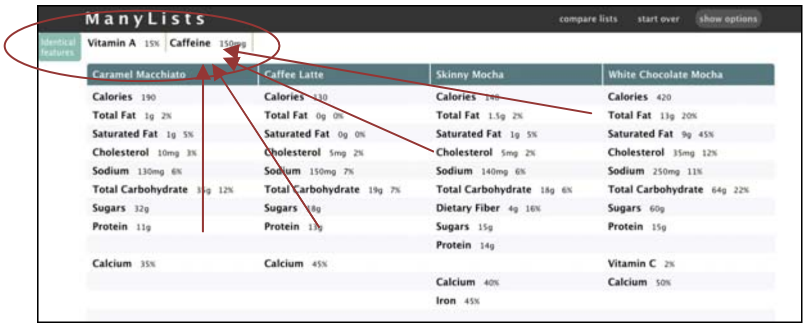

The identical features move to the horizontal strip at the top. All Vitamin A merge first, then all Caffeine, so we can see that all drinks have the same amount of vitamin A and caffeine.

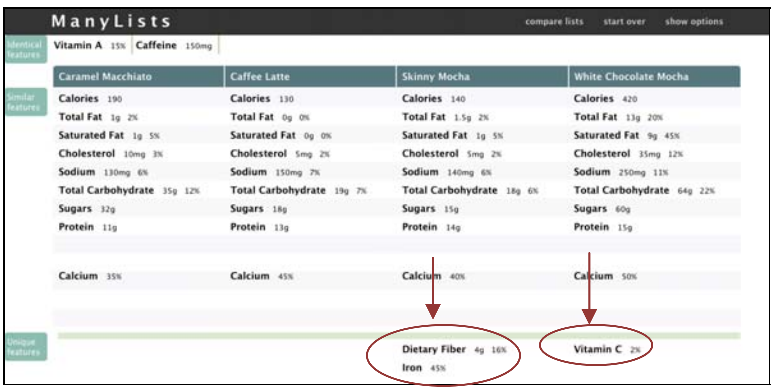

Features that are unique drop down to the bottom of each column, in the newly added "unique features" lower area, showing for example that only the 4th drink has vitamin C.

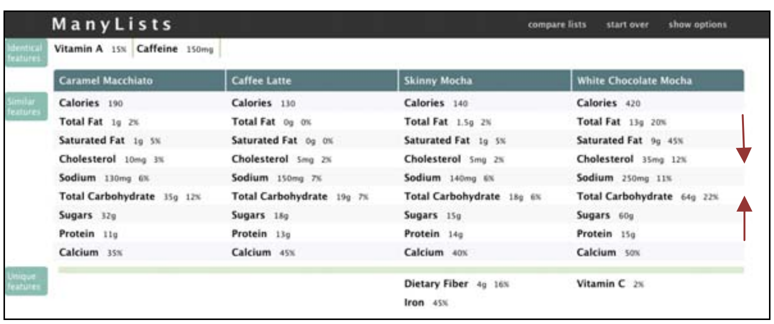

The identical features are aligned across all four columns

The best values are highlighted in green (since there is "goodness" associated with each feature, e.g. less fat is "good".

Figure 4: ManyLists explains the layout comparing four coffee drinks with multi-steps animation

The results suggest that the majority of participants could learn ManyLists' spatial groupings very easily. They found the step-by-step animation to be helpful in understanding the differences between products. The participant who needed help understanding the animation primarily had problems with the identical matching step. He reported that too many things were moving too fast and that he had not noticed the label for “identical features”. Similarly, a few other participants also complained about the animation for identical features, but they were able to discover the meaning of that spatial layout by reading the label. We believe that making the labels appear first, then later animating the items would be more beneficial. Highlighting the four copies of identical features before moving them to the top may also draw users' attention so they better perceive the initial location in the four columns. Another suggestion was to adjust the animation speed for the identical features: slow down the first group of identical features, and speed up for the remaining ones. Doing this will give users more time to understand the animation, but just not be too time.

Other findings related to the product selection application and not to the use of animation, e.g. all participants took into account the total number of best (green) features; some participants started with the two products with the most highlighted features, while others tended to start by elimination (e.g. the ones with the more calories). Almost all participants first narrowed their choices down to two options. Suggestions included letting users be able to eliminate products.

Discussion: A usability study did not allow us to measure the benefit of animation in isolation but gave us a lot of feedback. The majority of users understood the completely unfamiliar spatial groupings immediately, and provided useful suggestions for further improvement of the design of the animation. We had already included a step to highlight the color legend before highlighting best features, but users asked us to also highlight the labels of sections before moving the items. Overall we heard a strong message to slow down the animation.

SUMMARY OF DESIGN PRINCIPLES AND LESSONS LEARNED

Selecting the steps of the animation

We believe that the strength of Twinlist and ManyList animation lies in its well-designed series of steps. We broke a complex overall transition into separate steps, corresponding either to the movement of homogeneous groups of items (e.g. all the identicals), or to global reorganization with a clear visual effect (“compact” operation). Some steps should be easily divided further (e.g. the unique drugs could be further grouped on intake unique and hospital unique). The first step of the animation seems particularly important and should be even slower than other steps, helping users engage slowly. At first we animated all the identical drugs at once, but users seemed overwhelmed by this first complex step so we split it further, animating one pair of drug at a time, which seemed superior.

Timing

We allocated approximately 1 second per step, with up to 750 ms. to properly punctuate separate steps. User feedback suggests that this is the upper limit and that slower steps might be best. Multi-step animations seem to benefit from a unifying narrative. Items from the same set animate in unison; the sets themselves animate sequentially. Each set waits for its turn: items from the first set fly to the center and converge, then items from the second set fly to the center and converge, and so on. An important lesson is that it takes time to choreograph a series of steps into a natural sequence of events.

Timing can direct user attention to lessen the burden of comprehension. Each animation takes 800 msec., follows an ease-in, ease-out trajectory, and is punctuated with a 600 msec. pause. The base duration deviates slightly from the 1 second norm [Robertson et al. 2002], shortened to accommodate the 600 msec. necessary to properly separate different stages. The ease-in, ease-out pacing -- based on traditional principles of animation [Lasseter 1987] -- provides aesthetic fluidity and eases object tracking [Dragicevic et al. 2011].

The overall timing might in turn influence the preferable number of steps: for example, the merging animation of the identical drugs can become very long if many drugs are identical. If this is the case the merging of later pairs should be accelerated so as not to impact the overall duration of the animation.

User controls

The application should allow users to replay the entire animation, or to animate one step at a time. This can be used by expert users to demonstrate the interface to colleagues that are new users, and be useful for new users on their own. Finally, some users will want to turn off the animation entirely when they become familiar with the groupings, if don’t want to lose precious seconds to animation.

Adequate spatial groupings

For the animation to be effective, the groupings need to be meaningful. While grouping has long been recognized as a basic screen design principle [Shneiderman and Plaisant, 2010], there is not much research in this area.

One possible explanation for the benefit of animation in the above applications is that the animation matches the type of movement human do when they organize items. It is likely that observing people group items by hand on a small table will be helpful to design useful on screen grouping (similar to card sorting).

While animation can be helpful, it cannot to be used in isolation. The use of unifying background colors for different groups and of course informative labels will complement and support the animation.

The location of the groups needs to make use of the screen space efficiently. In Twinlist there was only two lists and they included 20+ drugs each max so the identical drugs can fit well in the center column. In ManyLists, with more lists to place on the screen the center column could not carry the same meaning, so we moved the identical items in a special section at the top of the screen, which also gave more space to the columns of similar items. We also used a horizontal list for the identical features, which saved vertical space. Both Twinlist and ManyLists include a compacting step to reclaim empty space, highlighting the constant relationship between time and space: it takes time to animate the compacting step but space is saved.

Visual representations of relatedness from other domains are abundant and may be useful to guide the choice of spatial groupings. [Kang at el. 2008] applied a five-column spatial layout to visualize entity resolution in relational data [13]. Jigsaw [Stasko et al. 2008] places special emphasis on visualizing connections, coordinating various views of lists of people or places to encourage dynamic document exploration. Parallel Tag Clouds [Collins et al. 2009] use faint, weighted "stubs" to indicate related items, converting these "stubs" to complete lines only when the user focuses on a particular list item; sizing by rank or score also intuitively assigns visual salience based on relative importance. The Semantic Graph Visualizer [Andrews et al. 2009] encodes structural similarities and differences with color, juxtaposing source and reconciled graphs for natural pairwise comparisons.

DISCUSSION AND CONCLUSIONS

Using two example applications, we presented a new beneficial use of animation, i.e. to help users learn and understand the meaning of spatial groupings representing similarities and differences between lists of items. Twinlist helped users compare two separate lists of medications. Manylists helped users compare products. Other interface styles are possible for those applications but statistical results of evaluation #1 and user feedback in user study #2 suggest that animation helped participants learn the spatial groupings we chose for Twinlist and ManyLists. User feedback convinced us of the need to further slow down our animations, and to add additional clarifying steps.

When asked, the majority of study participants stated that would be likely to turn off the animation during prolonged use, while others stated that they will want to keep it always. For the majority of users who would turn off the animation, one might argue that a canned animated training video might be just as beneficial as using an animation integrated in the application, and possibly easier to develop. On the other hand, training videos can only use sample training data while animations can act on users' own data, i.e. the data they will need to act upon. Further studies could measure if integrated animation is indeed more effective than recorded videos. Animating users' own data is probably more useful for intermittent users who need to re-learn the interface every time they use it.

Other comparison applications could make use of the techniques we presented here. For example, comparing language use by political candidates is a hot topic at the time of elections. Tag clouds are popular but the spatial groupings of Twinlist (for 2 candidates) or ManyLists (for 3 or more candidates) may be more effective.

We hope that designers will continue to explore multi-step animation and to report success and failures so that our combine experienced will lead to a better understanding of the power and dangers of animation.

ACKNOWLEDGMENTS

This work is supported in part by Grant No. 10510592 for Patient-Centered Cognitive Support under the Strategic Health IT Advanced Research Projects Program (SHARP) from the Office of the National Coordinator for Health Information Technology. We particularly thank our Sharp collaborators Todd Johnson, Eliz Markovitz, Shitra Shriram and main physician partners Jorge Herskovic, Elmer Bernstam, Zach Hettinger and Seth Powsner. Thanks also go to students Leonardo Claudino, Sameh Khamis, Ran Liu, Ben London, and Jay Pujara for developing the original prototype as a class project [Claudino 2011].

Websites: www.cs.umd.edu/hcil/sharp and www.cs.umd.edu/hcil/manylists

REFERENCES

1. Andrews, K., Wohlfahrt, M., Wurzinger, G. Visual Graph Comparison, *Proc. IEEE Symposium on Information Visualization* (2009) 62-67.
2. Bederson, B. B. and Boltman, A., Does animation help users build mental maps of spatial information?, in *Proc. of the 1999 IEEE Symposium on Information Visualization*, 11, 98 (1999) 28-35.
3. Bozzo Silva, Bernstam, Markowitz, Johnson, Zhang and Herskovic, Automated medication reconciliation and complexity of care transitions, *Proc. AMIA 2011*, 1252–1260.
4. Chevalier, F., Dragicevic, P., Bezerianos, A., and J. D. Fekete. Using text animated transitions to support navigation in document histories. *Proc. CHI '10*, (2010) 683–692.
5. Claudino, L., Khamis, S., Liu, R., London, B., Pujara, J., Plaisant, C., Shneiderman, B., Facilitating Medication Reconciliation with Animation and Spatial layout, *Proc.Workshop on Interactive Healthcare Systems (WISH2011)* --- Copyright retained by authors ---.
6. Collins, C, Viegas, F. B. and M. Wattenberg, Parallel Tag Clouds to explore and analyze faceted text corpora,

Proceedings of the 2009 IEEE Symposium on Visual Analytics Science and Technology (2009) 91-98.

7. Dragicevic, P., Bezerianos, A, Javed, W., Elmqvist, N., and J. D. Fekete, Temporal distortion for animated transitions, *Proc. Annual Conference on Human Factors in Computing Systems* (2011) 2009–2018.
8. Fisher, D., Animation for Visualization: Opportunities and Drawbacks, in *Beautiful Visualization: Looking at Data Through the Eyes of Experts*, O'Reilly Media, Inc., (2010) 329-352.
9. Guilmaine, D., Viau, C., McGuffin, M. J., Hierarchically animated transitions in visualizations of tree structures Graph visualization / *Proceedings of the 2012 International Conference on Advanced Visual Interfaces* (2012) 514-521
10. Grosjean, J., Plaisant, C., Bederson, B., SpaceTree: Supporting Exploration in Large Node Link Tree, Design Evolution and Empirical Evaluation, *Proc. IEEE Symposium on Information Visualization* (2002) 57-64
11. Heer, J. and Robertson, G.. Animated transitions in statistical data graphics. *IEEE Transactions on Visualization and Computer Graphics*, 13, 6 (2007) 1240–1247.
12. Kang, H., Getoor, L., Shneiderman, B., Bilgic, M., and Licamele, L., Interactive entity resolution in relational data: A visual analytic tool and its evaluation., *IEEE Transactions on Visualization and Computer Graphics*, 14, 5 (2008) 999 –1014.
13. Lasseter, J. , Principles of traditional animation applied to 3D computer animation, *ACM SIGGRAPH Computer Graphics*, 21, 4 (1987) 35-44.
14. Lawrence, A.W.; Badre, A.M.; Stasko, J.T.; , Empirically evaluating the use of animations to teach algorithms, *Proc. of IEEE Symposium on Visual Languages* (1994) 48-54
15. Markowitz, E., Bernstam, E., Herskovic, J., Zhang, J., Shneiderman, B., Plaisant, C., Johnson, T., Medication Reconciliation: Work Domain Ontology, Prototype Development, and a Predictive Model, *Proc. AMIA Annual Symposium (2011*) 878- 887.
16. Mayer, R., Multimedia Learning, 2nd ed ition (2009)
17. Novick, D., Rhodes, J., Wert, W., The communicative functions of animation in user interfaces, *Proc. Int. Conf. on Design of Communication* (2011)1-8
18. Plaisant, C., Shneiderman, B., Show Me! Guidelines for Producing Recorded Demonstrations, *Proc. IEEE VL/HCC'05* (2005) 171-178
19. Poon E. G., Blumenfeld, B. et al., Design and implementation of an application and associated services to support interdisciplinary medication reconciliation efforts at an integrated healthcare delivery network, *Journal of the American Medical Informatics Association*, 13, 6 (2006) 581–592.
20. Poon EG. Medication reconciliation: whose job is it? *AORN J.*, 89, 6 (2009) 1180, 1122.
21. Pronovost, P., Medication reconciliation: a practical tool to reduce the risk of medication errors, *Journal of Critical Care*, 18, 4, (2003) 201-205
22. Robertson, G., Cameron, K., Czerwinski, M., and Robbins, D.. Polyarchy visualization: visualizing multiple intersecting hierarchies. In *Proc. conference on Human factors in computing systems* (2002) 423–430
23. Robertson, G., Fernandez, R., Fisher, D, Lee, B., and Stasko, J., Effectiveness of animation in trend visualization. *IEEE Transactions on Visualization and Computer Graphics*, 14, 6, (2008) 1325–1332
24. Robertson, G., Mackinlay, J. and Card, S., Cone trees: animated 3d visualizations of hierarchical information. In *Proceedings of CHI* '91 (1991) 189–194
25. Rosling, H., Stats that reshape your worldview, *TED talk - online video* (2006, retrieved Sept.2012) www.ted.com/talks/hans_rosling_shows_the_best_stats_you_ve_ever_seen.html
26. Shanmugasundaram, M., Irani, P., and Gutwin, C., Can smooth view transitions facilitate perceptual constancy in node-link diagrams? *Proc. Graphics Interface* (2007) 71–78
27. Shneiderman, B. and Plaisant,C., *Designing the User Interface* , 5th edition (2010)
28. Stasko, S., Görg, C., and Spence, R, Jigsaw: supporting investigative analysis through interactive visualization, *Information Visualization*, 7, 2 (2008) 118-132.
29. Thomas B. H. and Calder, P. Applying cartoon animation techniques to graphical user interfaces. *ACM TOCHI*, 8, 3 (2001) 198–222
30. Tversky, B., Morrison, J. B. and Betrancourt, M., Animation: can it facilitate?, *International Journal of Human-Computer Studies*, 57, 4 (2002) 247–262.
31. Tversky, B. Visualizing thought. *Topics in Cognitive Science*, 3, 3 (2011) 499–535
32. Tversky,, B., Heiser, J., Lozano, S.C, .MacKenzie, and Morrison, J. *Enriching animations*. Cambridge University Press, Cambridge (2007).
33. Ware, C., I*nformation Visualization: Perception for Design*: Third Edition, Morgan Kaufmann Publ., San Francisco, CA (2012).
34. Yee, K.-P., Fisher, D., Dhamija, R. and Hearst, M. Animated exploration of dynamic graphs with radial layout, *IEEE Symposium on Information Visualization*, (2001) 43-50.