EUROGRAPHICS Symposium on Sketch-Based Interfaces and Modeling (2010)
M. Alexa and E. Do (Editors)

Exploring Usability and Learnability of Mode Inferencing in Pen/Tablet Interfaces

M. Negulescu [†], J. Ruiz[‡] and E. Lank[§]

University of Waterloo

Abstract

The inferred mode protocol uses contextual reasoning and local mediators to eliminate the need to access specific modes to perform draw, select, move and delete operations in a sketch interface. In this paper, we describe an observational experiment to understand the learnability, user preference and frequency of use of mode inferencing in a sketch application. The experiment demonstrated that those participants instructed in the interface features liked the fluid transitions between modes. As well, interaction techniques were not self-revealing: Participants who were not instructed in interaction techniques took longer to learn about inferred mode features and were more negative about the interaction techniques. Over multiple sketching sessions, as users develop expertise with the system, we find that they combine inferred mode techniques to speed interaction, and frequently make use of scratch space on the display to retrain themselves and to tune their behaviors. Our results inform the design of sketch interface techniques that incorporate noncommand features.

Categories and Subject Descriptors (according to ACM CCS): Information Interfaces and Presentation [H.5.2]: Miscellaneous—

1. Introduction

To allow the stylus to accomplish multiple content-based tasks like inking, erasing and editing, sketch interfaces typically incorporate a set of interface states, or modes of interaction. Although various techniques seeking efficiency in mode switching have been examined [LHGL05, RL07, HGB*06] it has been shown that there is a measurable cost associated with modes [RBL08]. As the number of interface modes increases, there is a growing need to develop intelligent mode switching techniques that provide low cost access to different interface operations. The inferred mode protocol [SL03] attempts to minimize mode cost by combining draw, select and delete operations in a single mode using contextual information and local mediator buttons.

Figure 1 depicts the inferred mode protocol's interaction paradigm. To partially eliminate the need to switch modes in the interface, the inferred mode protocol examines the gesture drawn by considering where the user clicked (on a selection or somewhere else on the display) and the drawn path (short or long, closed or open). For example, as shown in Figure 1a, if a user draws a circle that contains an object, they may wish to draw a circle or to select the object. In this case, the interface supporting inferred mode inks a circle and displays a local button mediator labeled with "Select?". A user can then select content using lasso selection by pressing the select button or they can leave the ink on the display by ignoring the button. As shown at the bottom of Figure 1a, if no content is inside the circle gesture, there is no ambiguity, and the circle is interpreted as ink. Similarly, Figure 1b depicts click selection. If a user clicks (inks a short stroke) on another stroke, the object is selected. If the user clicks in whitespace, then either an ink dot is placed on the screen, or, if selections exist, everything is deselected and no dot is placed on the screen. Figure 1c shows delete versus shading. Finally, Figure 1d shows translation behavior. If a user performs pen-down on a selected object and drags, the object is translated. However, a pen down and drag anywhere else results in deselecting all objects and drawing the gesture.

† mnegules@cs.uwaterloo.ca
‡ jgruiz@cs.uwaterloo.ca
§ lank@cs.uwaterloo.ca

DOI: 10.2312/SBM/SBM10/087-094

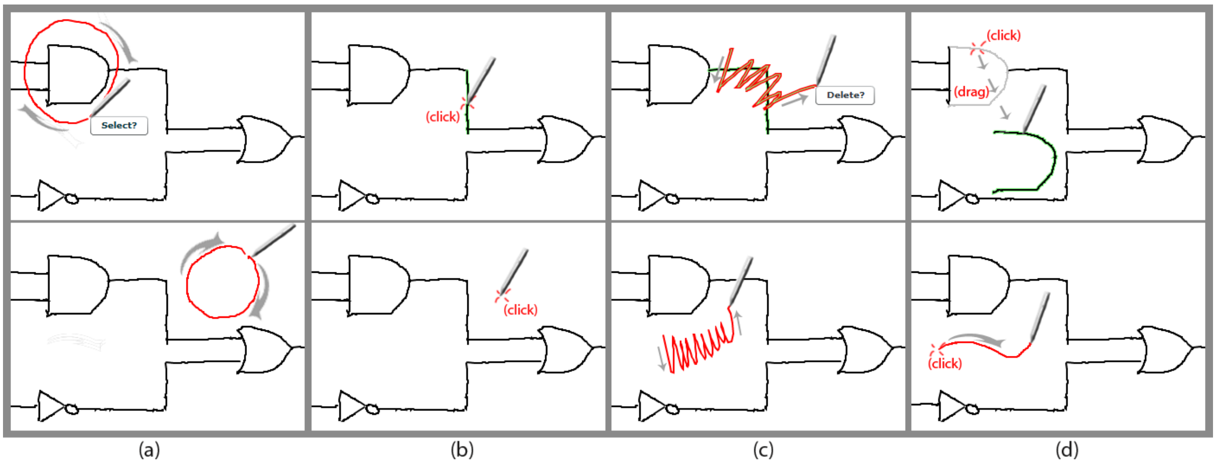

Figure 1: *The inferred mode protocol. Panel a. shows smart circle select. When an object is circled, a mediator appears (top), but no mediator appears if the circle encloses nothing. Panel b. shows smart select click. Panel c. shows smart delete (top) and shading (bottom). Panel d. shows translation (top) vs smart drawing (bottom).*

The goal of this research is to evaluate the inferred mode protocol as a tool for improving interaction in sketch-based interfaces. There are three specific aspects of the inferred mode protocol we explore. First, if the inferred mode protocol is available in an interface, alongside mode-based interaction, do users use the protocol, or do they just default to using interface modes? Second, how do users learn to use the protocol? Is it self-revealing? How does expertise develop over time? What new interactions do we see that take advantage of the protocol as users become more skilled? Finally, what are users subjective evaluations of the protocol and how might the protocol be improved?

Providing valid answers to these research questions requires careful experimental design and a longitudinal study [Nie00] of user behavior. To address this challenge, we designed a 2X2 observational study. We created two interface variants of the inferred mode, one in which the inferred mode was present by default and one in which the inferred mode had to be specifically invoked using a special "smart" mode; and two instruction conditions, one where participants were give a short overview of the inferred mode and one where no instruction was given. We analyzed screen videos of eight participants, two in each group, performing sketching tasks over multiple sessions each (over 30 one-hour sessions were analyzed). We conducted interviews of participants to capture impressions.

In our analysis, we examine issues of learnability and usability of the inferred mode technique. We find, first, that, without instruction, it is difficult to develop an accurate mental model of the interaction technique, and that participants rapidly become frustrated and ignore interface behaviors that they do not understand. However, if participants either understand (through instruction) or develop understanding (through exploration) of the technique, they use the technique liberally during their interaction with the sketch. Given an understanding of the technique, they also spend significant time exploring ways that the technique can be used to optimize their behavior within interfaces. For example, participants use scratch-space on the display to explore tolerances and improve their ability to invoke gestures; participants combined operations such as select and delete to perform group deletion in unused screen space; and participants treated mediators in different ways depending on the perceived cost of the operation (e.g. delete is more costly than select, even with an undo operation present).

While a full exploration of all implications of our data is outside the scope of an eight-page conference paper, we highlight some of the design implications of our research. We note the need for training on interaction techniques, the benefits of scratch space, and enhancements to inferred mode operations. Together, these results inform the design of new interaction techniques to support fluid inking and editing in interfaces.

2. Related Work

Many researchers have studied mode switching techniques in tablet interfaces. Researchers have

- Evaluated different mode switching techniques [LHGL05].
- Evaluated the cost of modes in interfaces with respect to time to switch modes and frequency of errors that modes cause [RBL08].
- Proposed various techniques (gestures [BZW*09, ZM06], making menus [OP09, HBRG05], handles [GBH09]) to partially or completely eliminate modes in interfaces.

The goal of our work is to evaluate the effectiveness of one technique that limits the onus on the user to switch modes, the inferred mode protocol [SL03]. Furthermore, our aim is to answer questions about learnability, expert behavior, and use over time of this technique in practical usage. In essence, our primary goal is to evaluate an intelligent sketch interface technique.

All of the techniques proposed to either simplify mode switching or to eliminate explicit modes from sketch interfaces have been evaluated experimentally. However, many of the evaluations performed have used simplified tasks like pie-cutting [LHGL05] or line drawing [RBL08]. Others have used discrete command invocation evaluation [HBRG05, GBH09] where the user is told to perform a specific command – 'delete', 'cut', 'copy', 'select' – and the user performs the action that invokes the command. While laboratory evaluations are useful in telling us about speed and error rate in controlled conditions, they tell us little about the usability of techniques in real-world drawing tasks [Nie00].

Some researchers, recognizing the shortcomings of pure laboratory experiments on controlled tasks, have performed studies with higher ecological validity in specific areas of sketch interface research. One example is Bragdon et al.'s [BZW*09] GestureBar evaluation, where participants performed diagram transcription and editing tasks that approximate real world diagram creation. Two gesture learning techniques, GestureBar and a crib sheet, were contrasted. However, as the goal of their study was restricted to evaluating a training mechanism for gestures, the participants were given a strict script/tutorial that described the tasks they must do in order to learn a particular set of gestures. The script was realistic, but was still carefully controlled and fully specified. Users had less agency than they would in self-directed tasks.

In user interface work, one example of a study with high ecological validity is Kurtenbach and Buxton's [KB94] analysis of marking menus in a real-world graphical user interface over a period of time. Two users used an application with marking menus for approximately ten hours in total, with use spread over several days. Although the evaluation had only two participants, the researchers were able to extract very personal continuous feedback that looked at both short-term usability and effectiveness and long term perceptions of the application through multiple extended sessions. This work, exploring real-world use of a small number of users in detail, was invaluable in validating many of the laboratory findings associated with the speed, accuracy, and learnability of marking menus.

3. Methodology

Our goal in this study was to evaluate the usability of the inferred mode protocol of Saund and Lank [SL03] during realistic sketching. Our study technique is inspired by Kurtenbach and Buxton's ecological evaluation of marking menus [KB94] and Bragdon et al.'s [BZW*09] evaluation of GestureBar. We describe an experiment where, as in Bragdon's evaluation, participants were given pre-specified sketch entry and editing tasks to perform in an interface incorporating the inferred mode protocol, but, as with Kurtenbach and Buxton's evaluation, participants were not required to use the inferred mode protocol to complete the sketching tasks.

We wished to measure user adoption of the inferred mode protocol, both from the perspective of learnability – how easy it is to acquire expertise with the technique – and user preference – whether users actually make us of the interface. Over time, user preference can be measured by comparing the frequency of use of inferred mode features with the frequency of use of other options available in the interface. If participants use either inferred mode or alternatives more frequently, we can claim that there is a preference for one or the other. How participants make use of the technique provides us with details on how expertise is acquired.

3.1. Task

The task was the entry and editing of a set of simple digital logic circuits. Participants were given an initial digital logic circuit and asked to draw it in the interface. They were then asked to modify the digital logic circuit in specific ways, for example by inserting, deleting, or changing gates. While the "work" done by participants was not real, they were free to perform the tasks in any way they wished within the sketch interface. We gave them no direction on how to perform the tasks, only what tasks they were to perform in the sketch interface. Tasks were performed on a Toshiba Portege M200 tablet with the application maximized at a 1400x1050 resolution.

3.2. Experimental Design

We designed a 2X2 observational study that looked at learnability and user preference of the inferred mode protocol. To study learnability, we divided our participants into two groups, those who received instruction and those who did not. All participants received approximately 5 minutes of instruction in digital logic circuits. All but one participant had some knowledge of digital logic circuits or of formal logic. The participants in the Instructed group were also give a three minute overview of how the inferred mode protocol worked in the sketching interface they were using, while participants in the Not Instructed group were given no information on the inferred mode protocol. To limit bias, we were careful to show participants in the Instructed group both the inferred mode protocol and mechanisms for changing modes in the interface without using the inferred mode protocol, and did not express any preference for one technique over the other. This design allowed us to determine how easy it was to master the inferred mode protocol. Was instruction necessary to master the interface technique, or was the technique self-revealing to users? What strategies did users in different conditions take to master the technique?

To study user preference, we wanted to see whether participants made use of the inferred mode protocol over time. To do this, we designed two interface variants. The first interface variant, pictured in Figure 2a, contained four modes: draw, select, delete and smart. The draw mode performed inking in the interface. Select allowed content to be lassoed or clicked on for selection, and translation operations could be performed on selected content for editing. The delete button allowed users to delete entire strokes by drawing a gesture that intersected strokes that they wished to delete. Finally, the smart mode button implemented the inferred mode protocol.

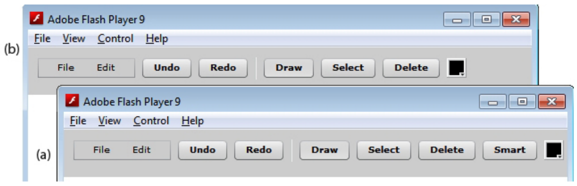

Figure 2: *The explicit interface in the foreground (a) and the implicit interface in the background (b). Note the extra "Smart" button in the explicit interface.*

When designing the study, one concern we had was that participants using the four-mode interface might never make use of the "Smart Mode" and, therefore, might never see any of the interface techniques that comprise the inferred mode technique. Participants were free to perform the tasks however they wished, and we wanted to ensure that at least some of the participants in our study saw the button mediators that invoke computational support. With this in mind, we designed a second interface.

The second interface variant (Figure 2b) had only three modes – draw, select and delete. Select and delete functioned identically to select and delete in the first interface. No computational intelligence was integrated into these modes. However, the Draw mode was designed to implement the inferred mode protocol, essentially mimicking the behavior of the "Smart" mode in the first interface.

In our study design, we label these interface conditions *explicit*, having an explicit smart mode, and *implicit*, having the smart mode implicitly included in draw mode. As a result of the two instruction and two interface designs, we had four unique configurations for our study: Instructed/Explicit, Instructed/Implicit, Not Instructed/Explicit, and Not Instructed/Implicit.

The study was designed as a between subjects, multi-session observational study. Each participant was assigned to one of the instruction/interface configurations, and remained with that instruction/interface configuration throughout their session (i.e. we did not use repeated measures). For each session, participants came to our lab and were given a set of drawing and editing tasks to perform, specifically a set of digital logic circuits to draw and then edit. Each drawing and editing session took approximately 45 minutes, and participants were paid $5 for each session they completed. Each participant completed between three and five sessions, as indicated in Table 1.

3.3. Observations

During each session, handwritten notes supported by screen capture videos recorded strategies and behaviors of participants. The videos were used to quantify the number of times button modes, inferred modes and mode errors occurred. After the last session, participants were interviewed in a semi-structured format for impressions and suggestions for improvements.

3.4. Participants

A total of eight participants (university students aged 21-26, 4 females) completed our study, two participants per condition. All of our participants were comfortable with computers and general graphics applications such as Paint or Photoshop, but had almost no experience with tablet interaction. One participant, P5, had some knowledge about pen/stylus interaction from watching YouTube videos on the topic.

While the small number of participants may give some pause, it is not unusual to perform qualitative studies designed for high ecological validity on a small sample. Consider, for example, Kurtenbach and Buxton's [KB94] work on marking menus, where only two participants were studied.

4. Results

In this section, we explore first the use of the inferred mode protocol during the drawing task by analyzing video data. We also look at the use of specific interaction techniques (circle select, click select, and scatch-out). Next we explore how participants learn to use the basic features of the protocol, examine training and retraining behaviors, and describe expert behaviors that evolved during our study. Finally, we present the subjective impressions of our participants from our post-experiment interviews.

4.1. Use of the Inferred Mode Protocol

The analysis of the use centers around an analysis of the video data gathered throughout the sessions. Though much of our study data is contained in observations of behaviour and from interview data post-sessions, Table 1 shows a summary of use for informational purposes. The first row indicates the number of sessions for each participant. Participants with higher levels of frustration used the application

	Instructed/Explicit		Instructed/Implicit		Not Instructed/Explicit		Not Instructed/Implicit	
	P1	P2	P3	P4	P5	P6	P7	P8
Number of Sessions	3	4	3	4	3	5	5	5
Smart Select Click	1.3	0.4	1.7	0.2	0.7	0.4	0.0	0.0
Smart Select Circle	11.7	4.6	11.7	8.5	3.7	0.2	0.0	0.0
Smart Delete	4.3	1.0	5.3	0.5	2.0	0.4	0.0	0.0
Smart Select Click Error	2.7	2.4	4.3	5.0	3.3	1.2	8.2	7.2
Smart Select Circle Error	0.3	1.0	1.3	0.2	0.0	0.0	0.0	0.0
Smart Delete Error	6.3	2.2	3.0	0.2	3.0	0.0	0.0	0.0
Ignored Select Circle	10.7	4.8	5.0	10.5	4.7	8.2	13.8	25.6
Ignored Delete	2.0	2.6	8.0	4.2	2.7	2.4	7.0	6.0
Button Select	0.0	3.0	0.0	0.0	2.0	6.0	9.4	8.6
Button Delete	8.0	5.8	0.7	7.8	4.7	7.2	7.4	9.0

Table 1: *Average frequency (number of uses per session) of operations, rounded to one decimal place. P1-8 represents our labeling of participants.*

for more sessions (up to five sessions) than participants with lower frustration, as it was our desire to evaluate inferred mode protocol's learnability and adoption. Therefore, participants who were least frustrated – P1, P2, P3, P4, and P5 – have fewer than five sessions. P6, P7, and P8 all have five sessions, an indication of their higher level of frustration.

The remainder of the table contains usage data on the inferred mode protocol and on other mode-switching techniques incorporated into the interface. The second grouping of data indicates the average number of times per session each participant used the inferred mode protocol interactions (labeled as *Smart Select Click*, *Smart Select Circle*, and *Smart Delete*). The third data grouping indicates the number of times participants tried to access inferred mode mediators and the mediators failed to appear or the inferred modes were activated in error. For example, a participant might draw a circular shape around an object that does not pass the threshold for recognition as a select circle operation. The mediator would, therefore, not appear, resulting in a *Smart Select Circle Error*. On the other hand, the participant might inadvertently click on an object, causing a selection action instead of a short pen stroke (a *Smart Select Click Error*). The fourth group of results, the *ignored* smart features, indicates those instances where a mediator appeared and users did not interact with it. These are not errors in the inferred mode's behavior. The inferred mode always assumes inking, and if participants want augmented behavior, they must interact with a button mediator. Finally, the last group, *Button Select* and *Button Delete*, are instances where participants used the explicit modes of operation.

When we examine the learnability of the inferred mode, it is easy to see that the Not Instructed group were much less likely to use the inferred mode features. For the Instructed participants, we see in Table 1 that these participants almost always used the inferred mode features, whereas the Not Instructed participants almost always use Button modes.

4.2. Individual Features

Participant use of the protocol uncovered common themes relating to usability and usefulness of the features in the protocol. These include participants' perceptions of mediators and of the different inferred mode techniques.

One characteristic of the inferred mode is the frequent presence of mediators on the display. Mediators were perceived very differently by participants in the Instructed and Not Instructed conditions. Instructed participants typically ignored the mediators when they did not want to select or delete content. However, participants in the Not Instructed/Implicit condition found the mediators annoying as they interfered with interaction on the display.

Participants P1, P2, P3, P4, and P5 used Smart Select Circle frequently as shown in Table 1. For all of these participants, it was easily the preferred mode for selecting content on the screen. Participants who made use of this feature saw it in the same light as a keyboard accelerator, with one participant even comparing it to a shortcut:

> *[Select Circle] is like a shortcut to me. I don't have to go to the menu and then back to the graph. [...] Very good feature to keep. [P4]*

Select Click was one feature that offered some controversy with our participants. Most participants found its use limited, but opinions varied as to how frustrating it really was. Similar to Select Circle, participants in the Not Instructed conditions found this feature frustrating. The problem with Select Click is inadvertent activations, common when dotting 'i's' or inserting punctuation. For example, P8, who did not use the inferred mode features, noted that it frequently selected things while he/she was drawing. Participants in the Instructed conditions used this feature sparingly, either because they forgot it existed or because the circle select was sufficient for their particular task. However, instructed participants found inadvertent activations to be little problem when sketching in the interface.

Only one participant, P3, used the delete gesture frequently; other participants used the delete gesture sparingly or not at all, as shown in Table 1. In our observations of participants, one thing that we observed is that a delete gesture is rarely the most efficient for deleting content. The explicit delete operation is more precise, making it preferred when deleting a stroke from a drawing. Delete gestures are not sufficiently precise to avoid surrounding content.

4.3. Learning Behaviors

When examining our video data, we observed an evolution in user behaviors. We separate these evolutions into three different themes. First, for users instructed in the inferred mode, we explore training and retraining behavior during interaction. Next, for participants in the Not Instructed condition, we explore active experimentation versus passive experiences of these participants as they worked with the interface. Finally, for participants who understood the functioning of the inferred mode, we describe expert-level characteristics of these users that can lead to the design of new sketch interface techniques.

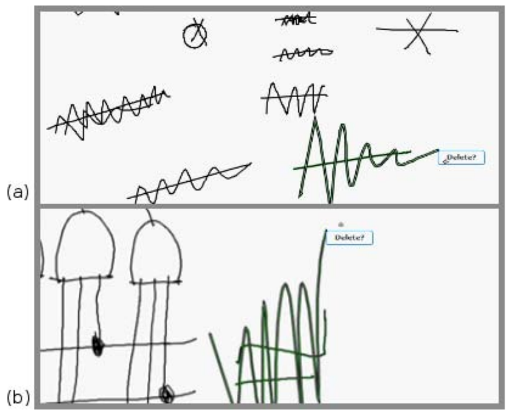

Figure 3: *A participant retraining in the delete gesture (a), and another scratching out accumulated strokes in their assigned "trash" region of the canvas (b).*

4.3.1. Training and Retraining

One common behavior of users who received instruction in the inferred mode was a training and retraining phase that we observed interspersed within the sessions. During the initial session, our instructed participants would try each of the features carefully to ensure that they could successfully activate the inferred mode. On subsequent sessions, participants would begin to perform their tasks. As they required features of the inferred mode, they moved to white space on the display and practiced the gestures prior to making use of the technique in their drawing. Figure 3a shows an example of one participant retraining the delete gesture.

This training and retraining behavior demonstrates the value of "scratch space" within sketch interfaces. In GestureBar [BZW*09] researchers noted the need for multiple interactions of a gesture to train muscle memory so the gesture could be replicated successfully on the display. As a result of this need for training or practice, Bragdon et al. created a separate panel activated from the toolbar for users to experiment on with gestures. In our interface, we also note the need for scratch space where users can explore the behaviors of the interaction techniques and the tolerances associated with mediator activation. While GestureBar accomplishes this using separate panels, our participants seemed to prefer interaction with content on the actual drawing canvas in an unused area and cleaning that area of the canvas up using a group delete operation after interaction. This distinction between participants in our study and participants in GestureBar's formative studies may be because of the distinction between a recognition interface like GestureBar, where each gesture must be recognized, and sketch interface techniques like the inferred mode protocol, where the user needs unrecognized ink content to interact with, i.e. to select or delete, in order to explore gesture behavior.

4.3.2. Active Experimentation

Participants in our Not Instructed condition (P5, P6, P7 and P8) were required to learn to use the inferred mode features without any explicit training, and we included no embedded help and no "What's This" widget to clarify the features, by design. Participants using an interface with embedded help would exhibit either the behavior of the Instructed participants – i.e. they would use the help feature to master the inferred mode techniques and then behave as Instructed participants – or the behavior of the Not Instructed participants. Our goal with our Not Instructed users was to explore initial encounters with the inferred mode to determine the strategies employed to learn to effectively use the technique.

For these participants, we noted an interesting and counter-intuitive contrast between participants with the Explicit interface, i.e. a separate "Smart" mode, and those with the Implicit interface, i.e. where the inferred mode protocol was on by default. Our impression was that the always-on nature of the Implicit interface would encourage participants to master the inferred mode faster, so that mediators were useful rather than a distraction, whereas participants using the Explicit would have little impetus to explore the behavior of inferred mode. In fact, the opposite was true; the presence of a "Smart" mode seemed to encourage experimentation.

Early behavior of participants in the Not Instructed condition was similar when interacting with the inferred mode initially (either by default or through the "Smart" mode). Participants learned that clicking away from the mediator would cause it to vanish. However, participants with the Implicit interface would then continue to passively encounter and dismiss the mediator, and made no apparent effort to master the inferred mode techniques. As noted by P8

> *It was more annoying because I didn't know when the boxes came up and they got in the way [...] I never bothered to do it because there were other ways to do it. [...] When it occured I know what it did but I didn't figure out how. [P8]*

Both participants also had five sessions to master the protocol, and, as we see from Table 1, at no time did they make use of any of the smart features despite ample opportunity to discover the protocol. For example participant P7 saw the Select mediator nearly 14 times per session, or 69 times over five sessions and P8 saw the same mediator 128 times, or almost twice as much. Despite the presence of these mediators, participants never explored their behaviors.

In contrast, participants in the Not Instructed/Explicit condition behaved quite differently. The existence of the "Smart" button prompted the curiosity of the two users in this condition, P5 and P6. P5 analyzed the Smart mode very carefully during the first session and mastered all of the smart mode techniques, then converged on the behavior of the instructed participants for later sessions. P6 also tried to master the "Smart" features during the first session but failed initially. However, during the fifth session this participant began to understand and use the features. As we limited our study to five sessions, we were unable to determine whether additional sessions would have increased P6's use of the Smart mode, and whether P6's behaviors would have converged on the behavior of our instructed participants.

4.3.3. Expert-Level Patterns

Extended use over multiple sessions leads to a familiarity with both the benefits and limitations of available tools. Participants who understood the functionality of the inferred mode (P1, P2, P3, P4 and P5) quickly began to combine operations to more effectively make use of the interaction techniques. First, the delete operation requires a scribble gesture made over content to be deleted. However, when content is densely arranged on the display, a scribble lacks precision to delete without affecting surrounding content. As a result, participants would move small content into white space and then delete away from surrounding content. Participants also used circle select in the inferred mode to accumulate strokes for later deletion. They would move the strokes to a "trash" location on the display, as shown in Figure 3b, and then delete all of the unwanted strokes at once.

As participants developed experience with the inferred mode techniques, we also noted some differences between their perception of delete and select operations. The delete operation removes content from the display, whereas the select content simply changes the state of on-screen content. The delete operation, even in the presence of "Undo", was treated more carefully by our users. We noted that some participants were very careful to dismiss the delete mediator to avoid accidental activation, while they were more comfortable simply ignoring the select mediator.

Finally, during post-experiment interviews, one participant, P2, noted that the inferred mode protocol was most valuable during editing operations. P2 used the Explicit interface, and noted that he/she would create initial content in draw mode and then move to smart mode to edit content, essentially treating the smart mode as a more convenient edit mode with edit and drawing operations seamlessly integrated. This observation of phases of work and the utility of the inferred mode during editing raises the possibility of plan-recognition research to tune sketch interfaces to specific tasks. If we can identify when users are drawing versus when they are editing content, we can explore automatically tuning the interface to their current task. Whether automated techniques for changing interface behaviors are more effective than an explicit mode the user enters to edit content is an open question, but one worth exploring.

4.4. Design Enhancements

During our post-study interviews, many participants suggested enhancements to the system. The most common recommendation was the implementation of a Delete Selection option. We saw participants create their own Delete Selection by selecting and moving objects from that region to an unused area on the canvas. The Delete Selection option would eliminate the translation operation.

We are experimenting with options for Delete Selection. One that appears to hold promise and maintains our default pen-and-paper behavior is a "select-then-cross" operation where users first select an object (using smart circle select) then draw a line through the selection to prompt to Delete. If they press the mediator, a delete occurs. Otherwise, the behavior defaults to pen-and-paper inking, and a line is drawn on the display and content is de-selected.

A second design suggestion involved options for eliminating click select in the inferred mode. P3 noted that selection and cutting of curves is a common and often tedious operation. Users first cut the curves. Then, if they deselect the objects, or if they drop the objects at another location and add to the end of the objects, it can become difficult to know where one stroke ends and the next begins. This participant felt that recovering selections would be simplified if there were a selection undo stack. Because much of the use of the select click feature is restricted to retrieving selections, an undo stack would eliminate the need for select click.

5. Discussion and Future Work

The inferred mode protocol is an example of Nielsen's non-command interaction paradigm [Nie93]. The premise of the inferred mode is that the role of the computer in supporting interaction is to "interpret user actions and [to do] what it deems appropriate" [Nie93]. Nielsen claimed that this form of interaction would dominate new user interface paradigms.

However, adoption has been slow, and realistic studies of interaction provide evidence for why this is the case.

When evaluating noncommand interaction in pen/tablet interfaces, we see many of the same pitfalls associated with past generations of intelligent interfaces [Tes81, Nor02]. For example, our users had difficulty developing mental models of how the inferred mode protocol worked. As we noted, the inferred mode protocol analyzes actions and context using a simple decision-tree model. Arguably, decision trees are the simplest form of computational intelligence, yet users still struggle to understand how the system works.

While it may seem that noncommand interfaces are difficult to understand, it should also be noted that many of our Instructed participants preferred the inferred mode protocol and used it extensively. Participants noted that changing from “Select” mode to “Draw” mode is much simpler with the inferred mode, as a user can start drawing at any location on the canvas. The challenge is in how best to communicate to participants the features that are available within intelligent interfaces like the inferred mode interface.

With this in mind, we note that creating an explicit mode for noncommand interaction seemed to work well. For our Not Instructed/Explicit participants, the “Smart” mode gave a clue that there was a non-standard aspect to the interaction, and motivated them to understand how the interface worked. Both participants actively explored the noncommand features of the inferred mode. P5 mastered the technique during the first session and behaved as an instructed participant, and P6 developed an understanding of the smart mode over five sessions of sketching. In contrast, making computational intelligence standard in the interface by embedding it directly into the “Draw” mode caused significant problems for our participants in the Not Instructed group.

In our study, we did not use context-sensitive help because we wished to create two distinct groups of participants – those who knew about inferred features, and those who had to learn on their own. In a real-world implementation of our system, likely all instruction would be through context-sensitive help, so a natural next step is to study how help on demand can be used to train users in smart sketch techniques.

6. Conclusions

In this paper, we explore the learnability and use of features of the inferred mode protocol using a multi-session observational study. We show that, with instruction, participants value intelligent interface techniques, and make liberal use of them during drawing. We also highlight lessons learned for incorporating noncommand behaviors into sketch interfaces in realistic settings.

7. Acknowledgements

Funding for this research was provided by the Natural Science and Engineering Research Council of Canada (NSERC), the Networks of Centres of Excellence for Graphics, Animation, and New Media (NCE-GRAND), and the Palo Alto Research Center (PARC, inc.).

References

[BZW*09] BRAGDON A., ZELEZNIK R., WILLIAMSON B., MILLER T., LAVIOLA JR. J. J.: Gesturebar: improving the approachability of gesture-based interfaces. In *CHI '09: Proceedings of the 27th international conference on Human factors in computing systems* (2009), ACM, pp. 2269–2278.

[GBH09] GROSSMAN T., BAUDISCH P., HINCKLEY K.: Handle flags: efficient and flexible selections for inking applications. In *GI '09: Proceedings of Graphics Interface 2009* (Toronto, Ont., Canada, Canada, 2009), Canadian Information Processing Society, pp. 167–174.

[HBRG05] HINCKLEY K., BAUDISCH P., RAMOS G., GUIMBRETIERE F.: Design and analysis of delimiters for selection-action pen gesture phrases in scriboli. In *CHI '05: Proceedings of the SIGCHI conference on Human factors in computing systems* (New York, NY, USA, 2005), ACM, pp. 451–460.

[HGB*06] HINCKLEY K., GUIMBRETIERE F., BAUDISCH P., SARIN R., AGRAWALA M., CUTRELL E.: The springboard: multiple modes in one spring-loaded control. In *Proceedings of the Conference on Human Factors in Computing Systems, CHI 2006* (2006), pp. 181–190.

[KB94] KURTENBACH G., BUXTON W.: User learning and performance with marking menus. In *CHI '94: Proceedings of the SIGCHI conference on Human factors in computing systems* (1994), ACM, pp. 258–264.

[LHGL05] LI Y., HINCKLEY K., GUAN Z., LANDAY J.: Experimental analysis of mode switching techniques in pen-base user interfaces. In *Proceedings of the Conference on Human Factors in Computing Systems, CHI 2005* (2005), pp. 461 – 470.

[Nie93] NIELSEN J.: Noncommand user interfaces. *Commun. ACM 36*, 4 (1993), 83–99.

[Nie00] NIELSEN J.: Novice vs. expert users, February 2000. http://www.useit.com/alertbox/20000206.html.

[Nor02] NORMAN D. A.: *The Design of Everyday Things*. Basic Books, September 2002.

[OP09] OAKLEY I., PARK J.: Motion marking menus: An eyes-free approach to motion input for handheld devices. *Int. J. Hum.-Comput. Stud. 67*, 6 (2009), 515–532.

[RBL08] RUIZ J., BUNT A., LANK E.: A model of non-preferred hand mode switching. In *GI '08: Proceedings of graphics interface 2008* (2008), pp. 49–56.

[RL07] RUIZ J., LANK E.: A study of the scalability of non-preferred hand mode switching. In *Proceedings of International Conference On Multimodal Interfaces, ICMI 2007* (2007).

[SL03] SAUND E., LANK E.: Stylus input and editing without prior selection of mode. In *UIST '03: Proceedings of the 16th annual ACM symposium on User interface software and technology* (2003), ACM, pp. 213–216.

[Tes81] TESSLER L.: The smalltalk environment. *Byte* (1981), 90–147.

[ZM06] ZELEZNIK R., MILLER T.: Fluid inking: augmenting the medium of free-form inking with gestures. In *GI '06: Proceedings of Graphics Interface 2006* (Toronto, Ont., Canada, Canada, 2006), Canadian Information Processing Society, pp. 155–162.