

# Sketcholution: Interaction Histories for Sketching

Zhenpeng Zhao<sup>1</sup>, William Benjamin<sup>2</sup>, Niklas Elmqvist<sup>1,3\*</sup>, Karthik Ramani<sup>2</sup>

<sup>1</sup>*Department of Computer Science, University of Maryland, College Park, MD, USA;* <sup>2</sup>*School of Mechanical Engineering, Purdue University, West Lafayette, IN, USA;* <sup>3</sup>*College of Information Studies, University of Maryland, College Park, MD, USA*

---

## Abstract

We present Sketcholution, a method for automatically creating visual histories of hand-drawn sketches. Such visual histories are useful for a designer to reflect on a sketch, communicate ideas to others, and fork from or revert to an earlier point in the creative process. Our approach uses a bottom-up agglomerative clustering mechanism that groups adjacent frames based on their perceptual similarity while maintaining the causality of how a sketch was constructed. The resulting aggregation dendrogram can be cut at any level depending on available display space, and can be used to create a visual history consisting of either a comic strip of highlights, or a single annotated summary frame. We conducted a user study comparing the speed and accuracy of participants recovering causality in a sketch history using comic strips, summary frames, and simple animations. Although animations with interaction may seem better than static graphics, our results show that both comic strip and summary frame significantly outperform animation.

*Keywords:* sketching, interaction history, animation, summaries, user study

---

## 1. Introduction

*Sketching* is commonly defined as rapidly created freehand drawing that is not intended to yield finished work, but rather to visually explore ideas [1], and is a

---

\*Corresponding author. University of Maryland, College Park, 2117H Hornbake Building, South Wing, College Park, MD 20742, USA. Phone: +1 (301) 405-7414, Fax: +1 (301) 314-9245.  
*Email addresses:* zhaoz@umd.edu (Zhenpeng Zhao<sup>1</sup>), william@purdue.edu (William Benjamin<sup>2</sup>), elm@umd.edu (Niklas Elmqvist<sup>1,3</sup>), ramani@purdue.edu (Karthik Ramani<sup>2</sup>)

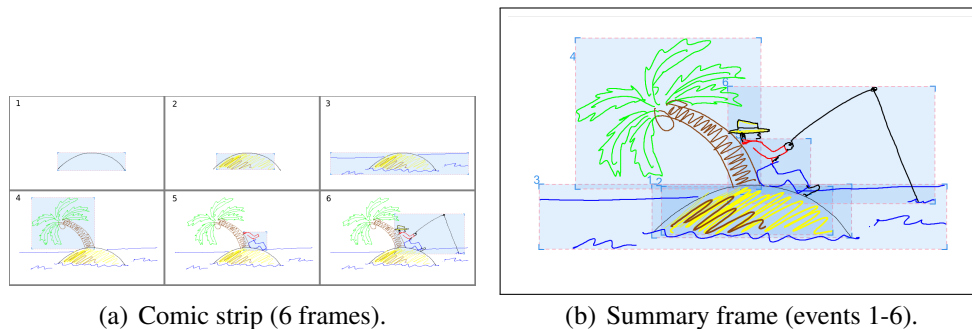


Figure 1: Our two Sketcholution visual history mechanisms illustrating the evolution of a sketch.

common tool for early design and creativity, regardless of discipline [2]. While pen and paper remains the most common medium for such activities, digital media for sketching has several compelling benefits beyond paper. One such is the ability to capture not just the final state of a sketch, but also every intermediate state along the way. Based on this idea, we propose *Sketcholution*, an automatic visual interaction history of how a sketch has evolved over time (Figure 1). Whereas a paper-based visual history would require the designer to take regular photocopies or digital pictures of the sketch being worked on, Sketcholution runs unobtrusively in the background of the digital sketch tool, capturing every single stroke made by the designer. This interaction history can then be played back, stroke by stroke, to show how the sketch was created and evolved over time. This would allow the designer to, for example, recall progress made during an earlier sketch session, communicate a particular idea to a collaborator or stakeholder, or access an earlier version of a sketch to either revert to or fork from that version (similar to source control systems).

Animation with interactive control is a seemingly obvious choice over static representation given that animations have been shown to improve understanding of spatiotemporal information [3, 4]. However, merely playing back an animation of the interaction history for a sketch is not necessarily the optimal presentation method. Complex animations can be difficult to perceive accurately [5] and are also potentially time-consuming to view in their entirety. For that reason, we propose two new *static* techniques for automatically summarizing sketch history captured during one or several design sessions: a *comic strip* of representative frames (or highlights) during the history, and a *summary frame* that annotates the changes made to the sketch in a single image. Both approaches rely on a bottom-up agglomerative clustering algorithm [6] that combines adjacent frames (each

representing a stroke) into *frame aggregates* while retaining the causal sequence of the interaction history. The decision of which frames to combine depends on the distance between consecutive frames as computed by a *frame distance function*. The resulting aggregation dendrogram can be cut at any level to yield a desired number of frames (for comic strips), or a particular distance threshold between events (for summary frames).

To determine which history presentation—animation, comic strip, or summary frame—is most efficient, we conducted a user study comparing the completion time of participants recovering the causal sequence of visual components in a sketch. Results show that both comic strip and summary are significantly faster than animation. These results also provide compelling evidence to the controversy surrounding animation for comprehension [7].

## 2. Related Work

Our sketch capture and summary mechanism lies at the intersection of sketching, early design, and interaction histories. Below we review relevant work in these research areas.

### 2.1. Sketching and Cognition

A *sketch* is a rapidly created freehand drawing that is not intended to create finished work, but rather to visually explore ideas [1, 2]. Sketches can be used in the design of electrical, mechanical, scientific, mathematical, and software artifacts [8]. For this reason, sketches are often used for idea generation and recording in early design—see below.

The order in which we sketch and draw reflects how we think. Taylor and Tversky [9] studied how people create regional maps and observed that the order in which people draw reflects their mental organization of the space. Regions which had features at multiple scales were depicted starting with larger features first followed by smaller ones. Sketches themselves have a definite structure similar to language and consists of basic elements such as lines and blobs [10]. Tversky [11] showed that the order of sketching elements reveals the designer’s underlying conceptual organization.

### 2.2. Sketching for Early Design

Design is a gradual, iterative process, often beginning from ill-defined or difficult problems that are decomposed, explored, and integrated in turns to yield many possible solutions [12]. Sketches play an important role in externalizing

ideas during early design [2], providing a “visible graphic memory” (p. 127) that facilitates creativity by providing an easily accessible database of generated ideas and by stimulating building on earlier ideas. Studies show that pictorial representations in general, and sketching in particular, is more effective than any other representations during early phases of ideation and creativity [13]. Furthermore, pen and paper remain the most common tools for sketching in early design [1]. Greenberg et al. [14] present a collection of methods to illustrate how to design with sketching.

Nevertheless, with the recent proliferation of pen-input devices, many efforts have been made to develop sketch-based interfaces for early design in a wide variety of domains such as architecture [15], automotive design [16], and software design [17]. These approaches all aim at replicating the good properties of paper—such as minimal learning curve, natural and precise interaction, and physical affordances—while retaining the unique benefits afforded by digital media, such as replication and composition of sketches.

### 2.3. *Interaction Logs and Graphical Histories*

Interaction logs and histories are common in human-computer interaction due to their relation to undo and redo operations, and modern user applications typically support multi-level versions of these, sometimes of a selective nature [18]. Similarly, navigation histories are central to web browsers, allowing users to easily go back and forward while browsing the Web. Heer et al. [19] carefully review the design space of interaction histories; we refer to their survey for further details on interaction capture and recall. Compared to Heer et al., our work targets a different media type and uses an aggregation algorithm to chunk the history.

Interaction data can be used to even greater effect. *Graphical histories* do not just maintain a list or stack of interactions, but also show them using a graphical summary [20]. This is most commonly done using a thumbnail image of previous state [19, 21], and allows for capturing interactions over time [22]. In recent work, Heer et al. [19] propose a comic strip-style graphical history using thumbnails of previous visualizations. Further, recent work has shown that augmenting sketch histories with contextual information such as pictures audio and videos improves the effectiveness of sketching for communication in early design [23]. Proper segmentation with users’ guidance makes the graphical history easier to understand [24].

Beyond interaction data, histories have also been used for summarizing other media types. For example, histories for binary image files can be modeled using a directed acyclic graph to store temporal and semantic relationships [25]. For video

histories, Barnes et al. [26] proposed continuous zooming to support navigation in time. Building on this, Ajmal et al. [27] give a comprehensive introduction to video summarization techniques, of which cluster-based and color-based methods are partially similar to our proposed aggregation approach. Eccles et al. [28] integrated geotemporal information into storytelling and presented stories with data such as behaviors and events. However, compared to all of the above techniques, our aggregation approach is different, our data type is sketch strokes, and we also focus on presentation techniques for the cluster data, which we evaluate with a user study.

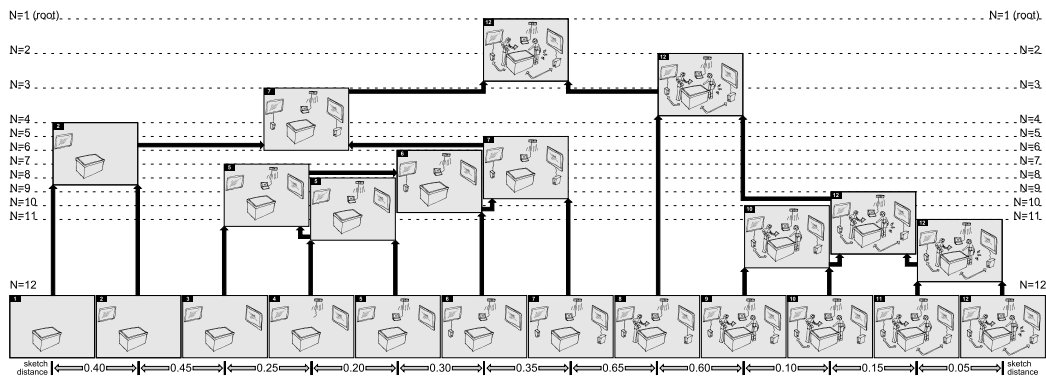


Figure 2: Visual dendrogram of sequential agglomerative clustering of a sample sketch. The distances between individual frames (bottom of image) only have to be computed once. Each sketch impostor shown in the dendrogram (above the bottom sequence) represents a cluster of two child sketches, and is chosen from the two children as the one with the most strokes (i.e., most information).

### 3. Sketcholution

*Sketcholution* is an automatic technique for capturing and visualizing the history of a sketch for the purpose of reflection, collaboration, and revision. Based on our literature review, we formulate the following design goals for a sketch-based history mechanism:

- **Stroke-level events:** A truly useful sketch history requires fine-grained history capture down to the level of individual strokes and formatting operations. This would enable the designer to reflect on each individual change in a sketch.

- **Chronology-preserving:** The causal order of how a sketch was constructed is a vital part of providing an accurate history. The chronology may also give insight into the creative process behind the sketch beyond the final result.
- **Efficient screen usage:** The sketch history should be space-efficient and adaptable to any screen area; in fact, it may become integrated into the sketch editor itself.
- **Efficient time usage:** An efficient history mechanism for sketches should support quick references so that it can become part of the sketch flow.

Many versioning systems allow users to revert and fork the version history from earlier revisions. However, such advanced version control operations are beyond the scope of this work; here we only concern ourselves with capturing and summarizing sketching.

### 3.1. Data Model

We define a *digital sketch* as a *canvas* and an ordered list of *sketch operations* that, when executed in sequence by an appropriate 2D rendering engine, yields a visual representation of the sketch on the canvas. Many types of sketch operations are potentially relevant; in this work, we include draw strokes, erase strokes, and formatting (color, transparency, stroke thickness). Both draw and erase strokes are modeled as pairs of 2D points (lines), where the former add digital ink to the canvas, and the latter remove it.

Using this definition of a digital sketch, it is clear that the sketch itself is also a stroke-by-stroke history of how it was constructed, from the first drawing operation to the last. We therefore call the list of sketch operations in the sketch a *sketch history*. This also means that rendering a sketch consisting of  $N$  operations from the beginning yields  $N$  consecutive *frames*—as in an animation—where the last frame  $F_N$  is the current state of the sketch:  $F_1, F_2, \dots, F_N$ . A frame  $F_t$  can be seen as the set of sketch operations added to the sketch up to a time  $t$ . To complete this interpretation, we also include time stamps  $t$  with each individual operation in the list, allowing us to exactly animate both the order and timing of the sketch.

### 3.2. Sketch Capture

Whereas most versioning systems (such as SVN, git, and Subversion) as well as cloud storage systems (Dropbox and Google Drive) operate on a file level,

Sketcholution must integrate with the sketch editor itself in order to collect stroke-level data. We focus on two main sketch operations—drawing and erasing—as well as changing the color and thickness of the drawing stroke, and the size of the eraser (however, other operations are also possible to capture and recall). All sketch input results in sequences of stroke segments (two points forming a line) being added to the sketch history as draw or erase strokes (depending on the chosen tool).

Storing a Sketcholution sketch amounts to serializing the canvas (dimensions, background, title, etc) as well as the entire sketch history. For long sketch histories, it may be impractical to execute or store the entire sequence of operations from beginning to end due to high rendering time and memory consumption, respectively. This is particularly true when many draw strokes have since been erased. For this purpose, we use the SVG format to store a cached version of the current state of a sketch’s visual representation.




Figure 3: Four-frame comic strip representation of the sketch shown in Figure 2. The four frames were chosen by cutting the aggregation tree (dendrogram) in Figure 2 at depth 3 ( $N = 4$ ), yielding four impostor children.

### 3.3. Sketch Summary

The design space of interaction histories is large [19], but our design goals limit the options to consider. Since our intention with this work is not to support forking past states, there is no branching model and thus no need for a hierarchical visual representation. Rather, the sketch history is a linear list, and the challenge instead becomes how to reduce the number of frames that are shown in the summary, from showing all  $N$  frames to just showing one.

In light of these constraints and goals, we formulate the following main types of sketch summary mechanisms (discussed below):

- **Animation:** Play back the entire sequence as a smooth animation. This option corresponds to one end of the above design spectrum: including all  $N$  frames.

- **Comic strip:** Render a subset of  $n$  representative frames as a static comic strip. The number would be selected as  $n \in (1, N)$ , so this constitutes a design compromise.
- **Summary frame:** Create a single frame that captures the entire sequence and enumerates the order of events. This design is at the other end of the spectrum:  $n = 1$ .

### 3.3.1. *Animated Playback*

Animated playback is the most straightforward of all: summarize the history of the sketch being drawn by simply replaying the operation sequence and show how the sketch evolves as strokes and formatting changes are executed. There is no attempt to reduce the number of frames, yet several design issues must still be resolved, the primary ones being timing, animation speed, and user control.

Timing refers to how to map the timestamped sketch operations in the sketch history to animated playback. While it may be useful to see the relative timing of operations (i.e., that the user took more time drawing one part of the sketch than another), most of the time the timestamps can be ignored and the operations simply be drawn in order with uniform timing. Speed, on the other hand, refers to the overall speed with which the animation is played back. Typically, an animated history will be played back faster than real time, although this could be left in the hands of the user. In addition, easing the speed of the animation in and out using temporal distortion may help perception [5]. Finally, user control concerns the interaction model that is used to control the animation; we use a media player metaphor, including play, pause, and rewind buttons, as well as a scrubber bar for directly controlling frame position.

### 3.3.2. *Sequential Aggregation*

Before being able to summarize a sketch using either a comic strip or a summary frame, we need to find a way to reduce the number of frames in the sequence, ideally without losing too much of the information encoded in the history. Furthermore, our design constraints mandate that we preserve the chronological order in the history, yet provide a flexible data structure that can adapt to available screen space. To achieve this, we use a variant of standard agglomerative (or hierarchical) clustering [6] where we preserve the order by only aggregating adjacent frames. This *sequential aggregation* algorithm preserves the chronology of frames, whereas traditional hierarchical clustering typically considers the distances between **all** items at each iteration of the algorithm.



For the aggregation to begin, we first need to compute the distances between each neighboring frame in the sketch history. The distance between two sketch frames  $d(F_{t-1}, F_t)$  is a measure of how much has changed between the two consecutive points in time, with the understanding that larger differences are more likely to be interesting than smaller ones. Several different distance metrics are possible (for example, how far the user’s pen has traveled from one frame to the next); we explore this in more detail later in this paper.

With the distances computed, the aggregation algorithm proceeds in the normal way by iteratively combining the two adjacent frames with the shortest distance into a single aggregate with the two frames as children. Distances for the new aggregate are taken from the child items, so no new distances need to be computed. The hierarchy created during this process is stored in memory, and the iterative merging is repeated until only one aggregate remains (the root of the aggregation hierarchy). Figure 2 shows how a sketch is iteratively clustered bottom-up into an aggregation hierarchy.

When creating a new aggregate, an impostor frame is chosen to represent the whole aggregate. Impostor selection should be stable, i.e., the impostor for an aggregate should be one of the two impostors of its children, otherwise the sketch history may change radically when expanding or collapsing an aggregate and thereby cause confusion. Of the two candidates taken from the children of an aggregate, our algorithm chooses the one with the most visible strokes because this branch of the aggregation hierarchy has the most information to convey to the user. Naturally, impostors for the original frames in the sketch history are those frames themselves.

### 3.3.3. *Comic Strip*

The basic idea behind the comic strip history mechanism is to select a subset of  $n \in (1, N)$  frames and show them in a static comic strip, i.e., in a list or grid where time runs from left to right, top to bottom. Figure 3 shows a four-frame comic strip of a sketch.

Of course, selecting *which* frames from the sketch history to include is the real challenge. This is where the sequential aggregation algorithm presented above comes in useful. Given a particular distance function, the dendrogram resulting from this cluster analysis constitutes a level-of-detail tree that can be cut at any level depending on the amount of screen space available. Such a cut (temporarily disregarding the hierarchy above the tree cut) yields a *forest* of  $n$  subtrees. The impostor frame for each subtree root becomes a representative frame to use in the comic strip. This way, the sketch history can be represented by any number of

frames  $n \leq N$ , including  $n = 1$  as well as  $n = N$ . For example, the comic strip in Figure 3 resulted from cutting Figure 2 at depth 3, yielding 4 subtrees represented by frames 2, 7, 8, and 12 from the sketch sequence.

#### 3.3.4. Summary Frame

Our intention with the summary frame was to provide a single frame where all events were captured. This would represent a compact, space-efficient visualization of a dynamic event sequence that is diametrically opposite from animation, which uses all frames. We achieve this by computing the union of all strokes (even those who were later erased) drawn at different points in time during the sketch history into a single frame, and then adding visual annotations to highlight and enumerate the major events during the history. Figure 4 shows an example of a summary frame for Figure 2.

Again we come back to the problem of identifying the major representative events during the sketch history, and the sequential aggregation algorithm again provides the solution. Let the *diameter* of an aggregate be defined as the sum of the sketch distances between all of its children. Instead of the user explicitly choosing the number of frames to display in a comic strip, we traverse the dendrogram depth first, only stopping the traversal for each branch when the current aggregate has a diameter less than or equal to a threshold value. The resulting list of frames become the events to highlight. For example, given a threshold of 0.45, such a traversal of the dendrogram in Figure 2 would yield frames 2, 5, 8, and 12. We then draw the bounding box of the strokes for each frame in the single summary frame and add an ordinal number. We use a green box and numeral for added strokes and red for erased strokes.

To be able to separate drawing and erasing into different events that can be conveyed with different colors, we must also use a distance metric that differentiates between these two types of events. This in turn will cause the aggregation hierarchy to become a forest instead of a single tree; a tree consisting of drawing strokes simply cannot be merged with a tree consisting of erasing strokes. A long sequence of alternate drawing and erasing will thus yield many events in the resulting summary frame, even if each is minor.

#### 3.4. Implementation

While we expect a production sketch tool to run on a mobile tablet, our prototype Sketcholution implementation (Figure 1) is built in Java using the Piccolo2D [29] library for vector graphics. The main component in the Sketcholution software is the sketch history, implemented as a queue of operations using

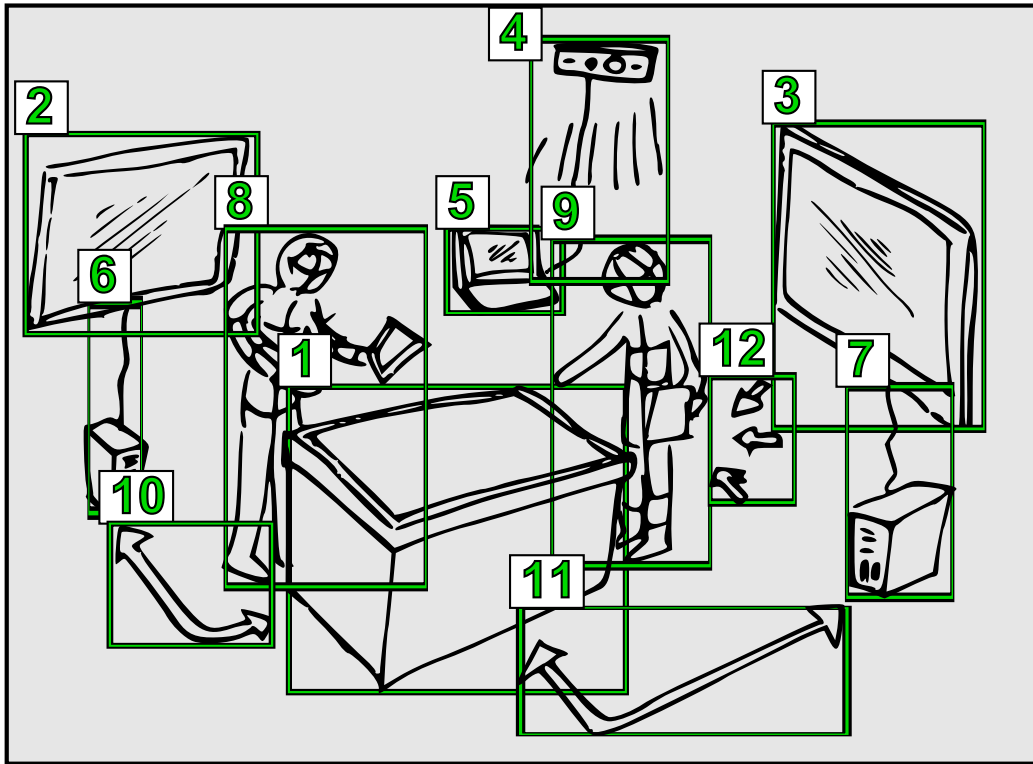


Figure 4: Summary frame for Figure 2 with ordinal numbers given the temporal sequence of events. The green bounding box communicates objects being added; removed objects (none in this sequence) would be drawn using a red color instead.

the Command design pattern. Because early design requires minimally-invasive interfaces and as close to natural interaction as possible, we chose not to expose this functionality in our prototype.

#### 4. Sketch Distance Metrics

Our sketch summarization approach has two major advantages compared to more traditional summarization approaches [30]: it is bottom-up, and it can be performed online. To achieve this, we need appropriate sketch distance metrics  $d(F_{t-1}, F_t)$  that accurately portray the absolute difference between two consecutive sketch frames at timestamps  $t-1$  and  $t$ . Below we review some of the distance metrics we have experimented with in our prototype:

- **Temporal distance:** The temporal order in how people sketch and draw

reflects how we think [9], which indicates that the time spent in drawing parts of a sketch may be an interesting distance metric to explore.

- **Euclidean distance:** Another simple metric is defined as the 2D distance (in pixels) that the user’s pen has moved between the frames. This metric will cluster adjacent strokes tighter than distant ones, and is used in Figure 1.
- **Perceptual distance:** Euclidean distance alone will not distinguish between color, thickness, and transparency, but such formatting changes are often important in segmenting a sketch history. A perceptual distance metric extends the Euclidean metric by assigning weights based on the visual saliency of the chosen stroke format.
- **Saliency:** A more complex distance metric would use the stroke saliency, where the saliency of a stroke at a specific point is related to the instantaneous curvature of the stroke. This local saliency can be found robustly by carrying out a local Principal Component Analysis. We do this by constructing a covariance matrix  $M$  of the coordinates within a small time window  $N(p_i) = \{p_j | j = i - w, j, i + w\}$  of width  $w$  and then performing an eigen-decomposition of  $M$  to yield eigenvectors  $(v_S, v_L)$  and eigenvalues  $(\lambda_S, \lambda_L)$  of  $M$ . We define the saliency at an instant as the ratio of the eigenvalues:

$$\alpha = \frac{\lambda_S}{\lambda_L} \tag{1}$$

For a perfectly straight line  $\lambda_S = 0$ , while for a circular region  $\lambda_S \cong \lambda_L$ . The distance between adjacent frames can be found by computing the  $\alpha$  difference:

$$d(F_{t-1}, F_t) = |\alpha_{t-1} - \alpha_t|. \tag{2}$$

This method is computationally efficient using a small frame window width  $w$ . From a pilot study (discussed below) we found that a frame width of 5 to 10 strokes is optimal to model intended sketch curvatures while filtering out small erratic moves.

Sketching is a highly subjective and application-dependent process. Hence it is difficult to come up with a single distance metric that fits all users and applications. In fact, composite distance metrics that weigh together more than one of the above

metrics are often good compromises. A general methodology to select distance metrics given an application is left for future work.

In many real-life design scenarios, such as that discussed by Schmidt et al. [31], continuous refinement and correction is an essential sketch activity. The final outcome has all the details and color on top of the original design scaffolding, which makes it hard for the user to review and have a quick overview for the whole process. In this case, a good distance metric may be based on color and thickness combined with temporal distance, as sketchers typically change their pens for either another color or thickness, and may often pause before starting another stage in the layered sketch.

## 5. User Study

At this point, we have derived three different techniques for conveying sketch history—animation, comic strip, and summary frame—but have no data on which of these techniques perform best under different conditions. Our intuition is that comic strips and summary frames will outperform animation when the testing sample is long and complex. On the other hand, for short and simple tasks, we expect that animation will be fastest since such tasks do not require summarization [27], and animation is also more familiar to participants.

To find the answer to this question, we performed a controlled user study that compared the completion time of participants recovering the temporal sequence of visual components being added and removed from a sample sketch.

### 5.1. Pilot Study

Prior to performing the full experiment, we first conducted an in-depth pilot study involving 16 participants. In the pilot study, we used a simple task sketch history recall task (similar to the actual study below), but used the pilot to determine suitable factors and conditions. During the pilot study, we found that the length of a sketch history and the context components that serve as background are interesting factors affecting the overall difficulty of the tests. However, the initial experimental design for the pilot study was not able to adequately distinguish between different difficulty levels. After increasing these two factors, we were able to obtain results so that groups were well separated based on completion time. We carefully selected the parameters to keep a balance so that participants will have enough challenge in the hard cases while their performance will not be affected by fatigue. This increased difficulty lead to increased session duration. In consequence, we were forced to take fewer participants in the formal user study due to

limited time and budget. We explain the details of the two experimental factors below.

For the techniques of Single Frame and Comic Strip, we also observed that obtaining the correct number of frames is essential to their effectiveness, i.e. the time for the participants to finish certain task using such techniques. Therefore, we developed an automatic segregation approach that separates the sketch into discrete frames. We tested this approach in the pilot study to reach an optimal separation rate, and used the parameters to the actual study with minor changes.

All of the parameters for factors in the below text were chosen to best separate the efforts of participants while keeping the tasks manageable.

### 5.2. *Participants*

We recruited 12 paid participants (6 male, 6 female) for this user study (each paid \$10). The participants were self-selected from the student population at our university, were aged between 20 and 31 years of age, had normal or corrected-to-normal vision, and were proficient computer users (all demographics were self-reported). A majority of participants were engineering students with a design background.

### 5.3. *Apparatus*

We conducted the experiment on a standard desktop computer equipped with a 24-inch LCD screen (resolution  $1600 \times 1200$ ), a standard keyboard, and a mouse. The Sketcholution software was maximized to fill the entire screen during the experiment.

### 5.4. *Task and Data*

Based on previous work on graphical histories [19] and causality in analytics [32], we selected a task based on recovering the causality of events in an interaction history. In our case, each trial consisted of the participant using one of the three sketch history mechanisms to view the evolution of a sample sketch.

The sample sketch consisted of several smaller “sketch objects”:  $D$  distractor objects,  $A$  added objects, and  $R$  removed objects. A sketch object is defined as a relatively small (approximately 100–300 pixels in dimension) sketched component of a clearly distinguishable physical item, such as a car, face, or heart shape. All objects were recorded using one-pixel wide black strokes for simplicity. A sample sketch was constructed by assembling a sequence of  $A$  sketch objects being added to the sketch, and  $R$  objects being removed. This design was chosen to allow us to procedurally generate new trials from scratch. Sketch objects were

selected from a pool of 23 discrete objects of different complexity, and each object was used at most once per trial. The order of object additions and removal was randomly permuted, whereas distractor objects never changed during a trial.

The user study task was to use the given sketch history mechanism to recover the order of added and removed sketch objects in the overall sketch. Participants were given a randomly ordered list of the sketch objects involved in the sample sketch, and were asked to specify the order. This was done by simply selecting a number representing its order of appearance or disappearance for each object using combo boxes (Figure 5). A trial could only be completed when the exact order for each sketch was given (incorrect clicks were recorded). The experimental platform silently recorded the completion time from the beginning of the trial to the moment the user clicked the submit button once the answer was correct.

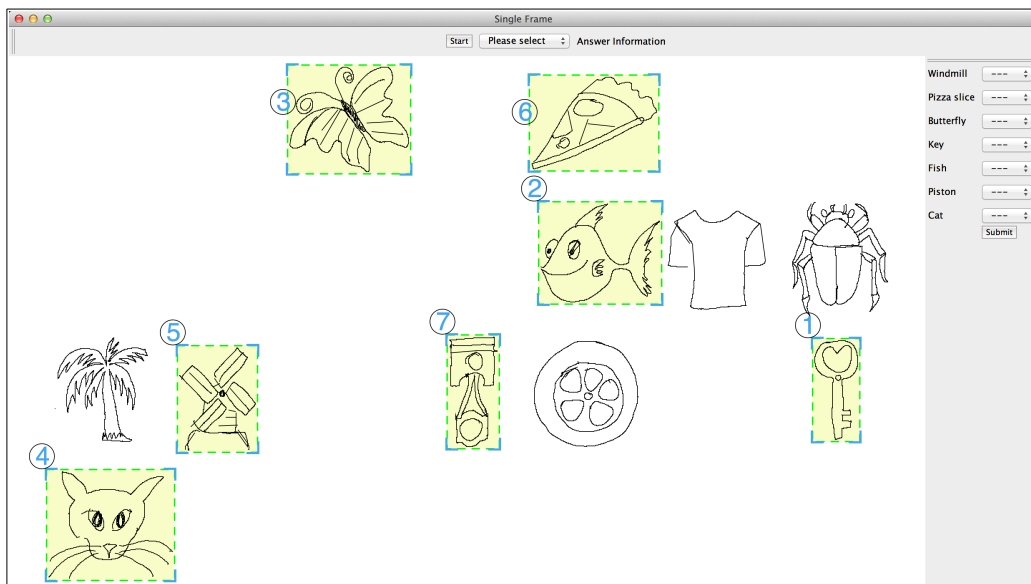


Figure 5: Screenshot from the study setup using the summary frame technique. In this example, we chose both temporal and Euclidean distance metrics for better segmentation. Temporal distance is added with a weight of 0.3 for better segmentation rates based on results from our pilot study.

## 5.5. Factors

We included the following factors in the experiment:

### 5.5.1. History Mechanism ( $H$ )

This factor modeled the sketch history mechanism  $H$  used:

- **Animation:** Basic animation playback where the participant can play, pause, and rewind the action, fast-forwarding, fast-rewinding, as well as moving the seeker bar to any time in the animation. All these controls were available throughout the trials. The playback window was stretched to fill the entire available screen space, minus the playback control interface.
- **Comic strip:** The Sketcholution comic strip where the available screen space is subdivided into smaller frames, each showing a representative snapshot of the history. The participant was able to control the number of frames using a slider; the initial number was five. Increasing the number caused the individual frames to become smaller, but showed a larger number of history highlights (and vice versa).
- **Summary frame:** The Sketcholution summary frame where a single annotated picture of the sketch history was stretched to fit the entire screen. The number of significant events to annotate in the summary was determined from distance thresholds found using a pilot. This number could also be controlled by the participant using a slider similar to the slider used for the comic strip; increasing the event number would simply highlight more events in the summary frame.

Even if our sketch objects were discrete, we did not artificially add delimitations between them, but used the *complete* stroke-level history of the sketch objects constituting the history.

### 5.5.2. Sketch Complexity ( $C$ )

The length of a sketch history will clearly have an impact on the difficulty of recovering the sequence of events; the longer the history, the more elements to survey and remember. To model this, we introduced a sketch complexity factor  $C$  as the sum  $C = A + R$ , i.e., the total number of added and removed sketch objects constituting the history. Given  $C$ , values for  $A$  and  $R$  were selected randomly. The pilot study had a low Sketch Complexity: 3, 5, 8. We increased it to 5, 8, 11 in the actual user study to increase the difficulty.

### 5.5.3. Distractor Level ( $D$ )

Sketches often contain components that either serve as context (for example, reference figures or background scenery), are part of earlier work, or were created by others in a collaborative setting. These visual components still add to the visual complexity of the sketch. We speculate that these invariant components in a sketch



that do not change over the course of the sketch history may therefore impact performance. We thus added a distractor factor  $D$  that modeled the number of such distractors that are visible in the sketch. We selected the distractor factor as a set of low (2 distractors), medium (4 distractors), and high (6 distractors) in the pilot study. We found that this configuration was not enough to serve its purpose, so we increased the number of distractors to 2 (low), 4–6 (medium), and 8–10 (high) distractors.

### 5.6. *Experimental Design*

We used a full-factorial within-participants design:

	12	participants
×	3	<b>History Mechanisms</b> $H$ (anim, strip, frame)
×	3	<b>Sketch Complexities</b> $C$ (5, 8, 11)
×	3	<b>Distractor Levels</b> $D$ (low, medium, high)
	324	Total trials (27 per participant, training excluded)

Trials were organized in blocks for each history mechanism. Block order was balanced using a Latin square participants to counteract learning effects; other factors were randomized. Trial completion time (in seconds) was the single dependent variable.

### 5.7. *Procedure*

An experimental session started with the participant arriving, reading and signing the consent form, and being assigned an identifier and history mechanism block order. The administrator then explained the general goals and task for the experiment. Each block started with the administrator demonstrating the use of the history mechanism and showing how to utilize it to solve a practice trial. The participant was then given two practice trials to solve. These trials were not timed and the participant was allowed to ask questions about the history mechanism and the task during this time. Since trials were timed individually, participants could rest between any two trials at will.

Individual trials started with an empty intermission screen that instructed the participant to click on a button to start the next trial. Clicking on this button caused the empty screen to be replaced by the history mechanism interface. A separate panel provided a randomly ordered list of the names of the sketch objects involved in the trial. The participant could select the ordinal number for each object's appearance or disappearance during the sketch history using a dropdown menu. There was no need to specify whether the object had been added or removed

during the sequence. Participants clicked on a submit button to finish the trial, but were allowed to proceed only if the provided order was correct.

After finishing a full block of  $3 \times 3$  conditions per mechanism, participants were asked to rate their perception of the efficiency, ease of use, and enjoyability of the history mechanism on a 1–5 Likert scale. After finishing all three blocks, they were asked to provide any general comments or feedback on any of the history mechanisms or the experiment. A full experimental session lasted approximately 60 minutes, including training and questionnaires.

### 5.8. Hypotheses

We formulate the following hypotheses about the study:

- H1 *Both Sketcholution interfaces will be significantly faster than animation.*  
We believe that the visual summaries provided by the comic strip and summary frame will be more efficient representations of sketch history than animation.
- H2 *High sketch complexities will cause significantly longer completion time than low complexities.* The number of discrete components being added or removed from the sketch will directly influence the time to solve a trial.
- H3 *Many distractors will cause significantly longer completion time than few distractors.* Analogously, many distractors will make the task more difficult, and thus slower.

## 6. Results

We here report on the completion times and subjective ratings for the three different history mechanisms.

### 6.1. Completion Time

We analyzed completion times using a repeated-measures analysis of variance (RM-ANOVA, all assumptions fulfilled) and found a significant main effect of history mechanism  $H$  on the completion time ( $F(2, 22) = 106.59, p < .0001$ ). Figure 6 shows a visual summary of the completion times; averages of 61.7 seconds for animation, 34.2 for summary frame, and 32.1 for comic strip. A posthoc Tukey HSD indicated that the difference between animation and comic strip as well as between animation and summary frame was significant ( $p < .05$ ), with animation on average having twice as long completion time as the other two mechanisms.

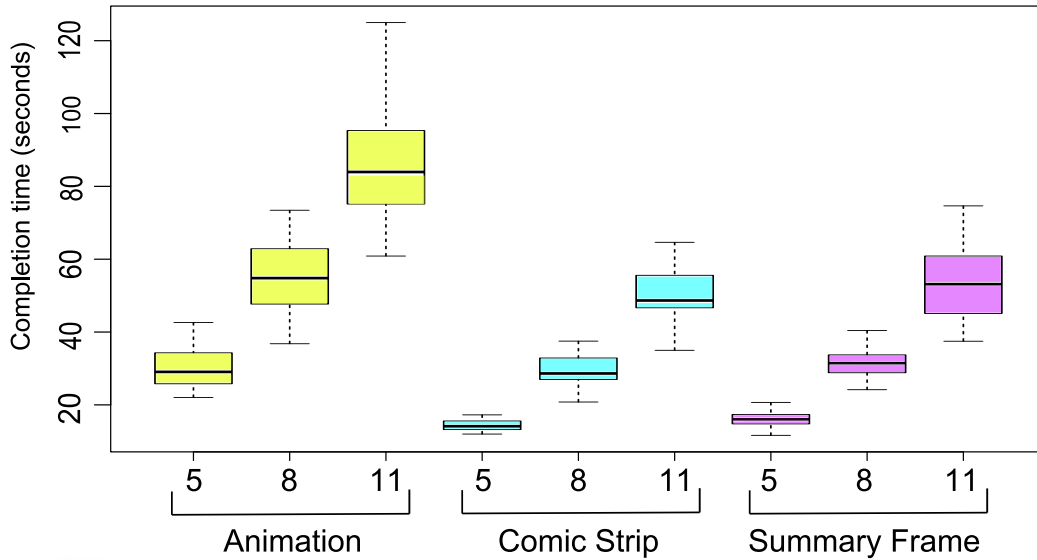


Figure 6: Average completion time as a function of history mechanism  $H$  and complexity  $C$ . Comic strip and summary frame show significantly faster times than animation.

There was no significant difference between completion times for comic strip and summary frame ( $p = .60$ ).

Not surprisingly, the sketch complexity  $C$  had a significant main effect on completion time ( $F(2, 22) = 215.67, p < .0001$ ). A posthoc Tukey HSD showed that all complexity levels had significant pairwise differences ( $p < .05$ ) in the order  $5 < 8 < 11$  in ascending completion time. Similarly, analysis of the distractor level  $D$  on completion time yielded a main significant effect ( $F(2, 22) = 4.89, p < .01$ ). A pairwise Tukey HSD showed significant differences for  $2 < 10$  number of distractors ( $p < .05$ ).

Finally, we found a significant interaction between history mechanism  $H$  and complexity  $C$  ( $F(4, 44) = 8.16, p < .0001$ ). A Tukey HSD posthoc test indicated that animation had comparable performance to both comic strip and summary frame at low complexity (5 events), but was significantly ( $p < .05$ ) slower for 8 and 11 events.

## 6.2. Subjective Ratings

Figure 7 depicts boxplots of the subjective ratings for each history mechanism on efficiency, ease of use, and enjoyability. We analyzed the 5-point Likert scale of subjective ratings for the different history mechanisms and found that

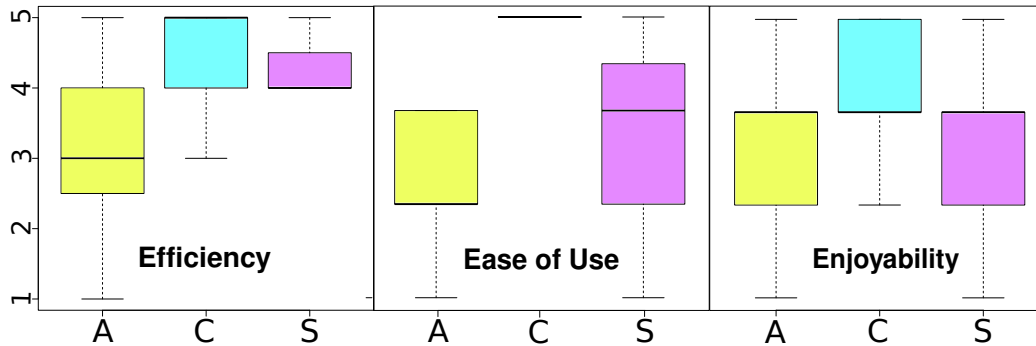


Figure 7: Subjective ratings (Likert 1–5). Comic strip and summary frame are higher than animation for efficiency and ease.

the efficiency and ease of use rating were significantly different across the three mechanisms (Friedman test,  $p < .05$ ), but the enjoyability had no significant difference (Friedman tests,  $p = .12$  for enjoyability). A Friedman posthoc test showed significant differences of efficiency between comic strip and animation ( $p < .05$ ), and significant differences of ease of use between comic strip and animation ( $p < .001$ ) and between comic strip and summary frame ( $p < .05$ ).

## 7. Discussion

We can summarize our findings as follows (discussed below):

- Both comic strip and summary frame are considerably faster than animation (accepting **H1**);
- The complexity of trials (i.e., the number of components) had a direct impact on trial difficulty (accepting **H2**); and
- The number of distractors in the sketch history had significant impact on completion time (accepting **H3**).

### 7.1. Explaining the Results

Our findings were in line with our initial hypotheses. Both summary frame and comic strip showed better performance than animation. The most likely explanation for these results is that our two novel methods allow the user to overview the entire event sequence without relying on a dynamically changing visual representation and requiring no user interaction. Users were simply faster in recovering causal order using Sketcholution for these reasons.

We made several interesting observations on how participants used the different techniques to solve the task. For comic strip, participants seemed to move from frame to frame using the bounding boxes to mark changes. The comic strip partitions events into an orderly grid, and so is faster than an animation. Similarly, the numbered bounding boxes in the summary frame imposes structure on the event sequence, relieving the user from having to replay the sketching history and explicitly remember the ordering.

It is hardly surprising that the complexity of a trial is indeed governed by the number of sketch components that are added or removed during the history, and this is a direct impact from the need to provide more information to the testing platform to recover the sequence. Animation of complex sketches involved a noticeably large amount of user interaction to recognize and recall events in the sketching sequence. The summarized nature of comic strip and summary frame made it significantly easier to identify the sketching sequence in the presence of complexity.

Distractors have a significant impact on completion times. This is again not surprising: a higher number of distracting items clearly makes it more difficult to memorize the event sequence, especially for animation. On the other hand, comic strip and summary frame have bounding boxes to indicate the changes which makes the distractors less effective. During the actual test, the combination of high sketch complexity and distractor level significantly increased the completion time.

## 7.2. *Generalizing the Results*

History mechanisms are pervasive in HCI, so our findings are potentially important beyond the sketch interactions discussed here. For example, it is notable that comic strips are common in several graphical history mechanisms [19]. Even though animation has comparable enjoyability to summary frame, it requires user interaction to start, stop, and rewind. In contrast, both summary frame and comic strips are static and require no user navigation.

On a more general level, both Sketcholution techniques proposed in this paper are examples of *spatialization* [33]: transforming temporal data into spatial representations. For Sketcholution, the stroke-level interaction sequence represents the temporal data, and the techniques only differ in which spatial representation is used. Comic strips use small multiples, whereas the summary frame instead orders the temporal data in the same image. Both approaches have their own pros and cons, and more research is needed to investigate the relative merits of each technique in detail.

However, the common denominator for all spatialization methods is that they run the risk of introducing high visual clutter and complexity. This is certainly also true for the Sketcholution techniques; the comic strip scales down and multiples the sketch canvas, yielding increased clutter proportional to the number of frames as well as smaller visual resolution for individual frames, and the numbering and bounding boxes used in the summary frame will cause similar levels of increased clutter and visual complexity. The saving grace for sketches in comparison to, for example, digital video is that the overall level of visual detail tends to be relatively low for sketches. However, it is equally true that a sketch will become increasingly complex as it is worked on over a longer period of time. Our Sketcholution techniques will not scale well with such long sequences and more work is needed to accommodate them.

There are several strategies that we could adopt to combat scale and complexity in the future. Simple optimization rules could be applied to the interaction sequence that eliminates redundancy, such as discarding strokes that were immediately deleted or undone. Multiple operations could be *chunked*, similar to our work in the skWiki system [34], where a series of consecutive stroke or erase operations are grouped together into a single meta-operation. Finally, as observed earlier in this paper as a rationale for introducing distractors in our user study, there are often parts of a sketch that do not change much. Such invariant components could be ghosted so that they are less visually salient and thus less of a distraction.

Another limitation of the Sketcholution approach is that sequential agglomerative clustering may not work well with holistic sketching where the illustrator evenly adds detail across the entire canvas (i.e., a form of breadth-first sketching) as opposed to working on different parts of the sketch separately (i.e., depth-first). We may have to relax the chronology-preserving design goal to better accommodate such behavior. We leave this for future work.

Finally, our intended use-case for Sketcholution is to embed the methods into the skWiki collaborative multimedia system [34]. skWiki stores persistent versions of sketches, rich text, and images using the concept of a path as a sequence of operations yielding a particular data object. We plan to use the Sketcholution method to visualize the evolution of a sketch in the system. While our study shows clear advantages to the comic strip and summary frame techniques proposed in this paper, a realistic implementation in the skWiki system may still provide all three options. In fact, hybrid solutions may also be desirable: for example, an animation could highlight recent changes using bounding boxes similar to the comics strip and summary frame. Another approach may be to chunk the animation into

individual segments, resulting in a hybrid animated comic strip.

## **8. Conclusions and Future Work**

We have presented Sketcholution, a new approach to capturing and summarizing the interaction history for digital sketching. The approach is based on an order-preserving agglomerative clustering algorithm that summarizes the stroke-level evolution of a sketch from the bottom up, yielding a visual summary that can be adapted to any desired level of detail. Comic strip and summary frame are two concrete history mechanisms that use the aggregation tree resulting from this cluster analysis, and they present the results as a sequence of highlighted frames, or as a single annotated frame, respectively. Results from a controlled user study indicate that both comic strip and summary frame significantly outperform animation. Furthermore, both are static and require no user interaction.

In the future, we will keep exploring applications in the domain of artistic sketching where our techniques can be integrated into a tool of segmenting sketches based on their time stamps. Our techniques can also be handy in the field of technical sketches and sketches of process schematics since such sketches tend to have multiple layers where our tools can separate the overlapped layers in certain cases. Our future work will also study how to use findings from the user study in a revision control framework for sketches. It should be possible to integrate our Sketcholution techniques directly into the editor interface. Furthermore, beyond vector drawing, automatically extracting animation summaries is appealing for many domains, including storyboarding, general interaction histories, and video summarization.

## **9. Acknowledgments**

This work is partly supported by the National Science Foundation grant no. 1227639, the Donald W. Feddersen Chaired Professorship, and the Purdue School of Mechanical Engineering. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsors.

## **References**

- [1] D. G. Ullman, S. Wood, D. Craig, The importance of drawing in the mechanical design process, *Computers and Graphics* 14 (2) (1990) 263–274.

- [2] R. H. McKim, *Experiences in visual thinking*, Brooks/Cole Pub. Co, Monterey, CA, 1972.
- [3] L. P. Rieber, Using computer animated graphics in science instruction with children, *Journal of Educational Psychology* 82 (1) (1990) 135.
- [4] B. Zhu, M. Iwata, R. Haraguchi, T. Ashihara, N. Umetani, T. Igarashi, K. Nakazawa, Sketch-based dynamic illustration of fluid systems, in: *ACM Transactions on Graphics*, Vol. 30, ACM, 2011, pp. 134:1–134:8.
- [5] P. Dragicevic, A. Bezerianos, W. Javed, N. Elmqvist, J.-D. Fekete, Temporal distortion for animated transitions, in: *Proceedings of the ACM Conference on Human Factors in Computing Systems*, Vancouver, BC, Canada, 2011, pp. 2009–2018.
- [6] A. K. Jain, M. N. Murty, P. J. Flynn, Data clustering: A review, *ACM Computing Surveys* 31 (3) (1999) 264–323.
- [7] B. Tversky, J. B. Morrison, M. Betrancourt, Animation: Can it facilitate?, *International Journal of Human-Computer Studies* 57 (4) (2002) 247–262.
- [8] I. E. Sutherland, Sketch pad a man-machine graphical communication system, in: *Proceedings of the SHARE Design Automation Workshop*, ACM, 1964, pp. 6.329–6.346.
- [9] H. Taylor, B. Tversky, Descriptions and depictions of environments, *Memory and Cognition* 20 (5) (1992) 483–496.
- [10] B. Tversky, J. Zacks, P. U. Lee, J. Heiser, Lines, blobs, crosses and arrows: Diagrammatic communication with schematic figures, in: *Proceedings of the International Conference on Theory and Application of Diagrams*, Springer Berlin Heidelberg, 2000, pp. 221–230.
- [11] B. Tversky, What does drawing reveal about thinking?, in: J. S. Gero & B. Tversky (Eds.), *Visual and Spatial Reasoning in Design*, Sydney, Australia, 1999, pp. 93–101.
- [12] E. Taborda, S. K. Chandrasegaran, K. Ramani, ME 444: Redesigning a toy design course, in: *Proceedings of the International Symposium on Tools and Methods of Competitive Engineering*, Karlsruhe, Germany, 2012, pp. 597–608.



- [13] F. L. Mckoy, N. Vargas-Hernández, J. D. Summers, J. J. Shah, Influence of design representation on effectiveness of idea generation, in: Proceedings of the ASME Design Engineering Technical Conferences and Computers and Information in Engineering Conference, Pittsburgh, PA, USA, 2001, pp. 1–10.
- [14] S. Greenberg, S. Carpendale, N. Marquardt, B. Buxton, Sketching User Experiences: The Workbook, Morgan Kaufmann Publishers Inc., 2011.
- [15] J. Dorsey, S. Xu, G. Smedresman, H. Rushmeier, L. McMillan, The mental canvas: A tool for conceptual architectural design and analysis, in: Proceedings of the IEEE Pacific Conference on Computer Graphics and Applications, 2007, pp. 201–210.
- [16] L. B. Kara, K. Shimada, Supporting early styling design of automobiles using sketch-based 3D shape construction, *Computer-Aided Design & Applications* 5 (6) (2008) 867—876.
- [17] Q. Chen, J. Grundy, J. Hosking, An e-whiteboard application to support early design-stage sketching of UML diagrams, in: Proceedings of the IEEE Symposium on Human Centric Computing Languages and Environments, 2003, pp. 219–226.
- [18] T. Berlage, A selective undo mechanism for graphical user interfaces based on command objects, *ACM Transactions on Computer-Human Interaction* 1 (3) (1994) 269–294.
- [19] J. Heer, J. D. Mackinlay, C. Stolte, M. Agrawala, Graphical histories for visualization: Supporting analysis, communication, and evaluation, *IEEE Transactions on Visualization and Computer Graphics* 14 (6) (2008) 1189–1196.
- [20] D. Kurlander, S. Feiner, Editable graphical histories, in: Proceedings of the IEEE Workshop on Visual Language, New Orleans, Louisiana, USA, 1988, pp. 127–134.
- [21] K.-L. Ma, Image graphs — A novel approach to visual data exploration, in: Proceedings of the IEEE Conference on Visualization, San Francisco, California, USA, 1999, pp. 81–88.

- [22] M. Nancel, A. Cockburn, Causality: A conceptual model of interaction history, in: Proceedings of the ACM Conference on Human Factors in Computing Systems, Toronto, Ontario, Canada, 2014, pp. 1777–1786.
- [23] G. Li, X. Cao, S. Paolantonio, F. Tian, SketchComm: A tool to support rich and flexible asynchronous communication of early design, in: Proceedings of the ACM Conference on Computer Supported Cooperative Work, Seattle, Washington, USA, 2012, pp. 359–368.
- [24] G. Noris, D. Skora, A. Shamir, S. Coros, B. Whited, M. Simmons, A. Horning, M. Gross, R. Sumner, Smart scribbles for sketch segmentation, *Computer Graphics Forum* 31 (2012) 2516–2527.
- [25] H.-T. Chen, L.-Y. Wei, C.-F. Chang, Nonlinear revision control for images, *ACM Transactions on Graphics* 30 (2011) 105:1–105:10.
- [26] C. Barnes, D. B. Goldman, E. Shechtman, A. Finkelstein, Video tapestries with continuous temporal zoom, *ACM Transactions on Graphics* 29 (2010) 89:1–89:9.
- [27] M. Ajmal, M. Ashraf, M. Shakir, Y. Abbas, F. Shah, Video summarization: Techniques and classification, in: *Computer Vision and Graphics*, Vol. 7594, Springer Berlin Heidelberg, 2012, pp. 1–13.
- [28] R. Eccles, T. Kapler, R. Harper, W. Wright, Stories in geotime, *Information Visualization* 7 (1) (2008) 3–17.
- [29] B. B. Bederson, J. Grosjean, J. Meyer, Toolkit design for interactive structured graphics, *IEEE Transactions on Software Engineering* 30 (8) (2004) 535–546.
- [30] R. C. Gonzales, R. E. Woods, *Digital Image Processing*, Pearson Prentice Hall, Upper Saddle River, NJ, 2008.
- [31] R. Schmidt, T. Isenberg, P. Jepp, K. Singh, B. Wyvill, Sketching, scaffolding, and inking: A visual history for interactive 3D modeling, in: Proceedings of the ACM Symposium on Non-photorealistic Animation and Rendering, San Diego, California, USA, 2007, pp. 23–32.
- [32] B. chul Kwon, W. Javed, S. Ghani, N. Elmqvist, J. S. Yi, D. Ebert, Evaluating the role of time in investigative analysis of document collections, *IEEE*

Transactions on Visualization and Computer Graphics 18 (11) (2012) 1992–2004.

- [33] W. Javed, N. Elmqvist, ExPlates: Spatializing interactive analysis to scaffold visual exploration, *Computer Graphics Forum* 32 (2) (2013) 441–450.
- [34] Z. Zhao, S. K. Badam, S. Chandrasegaran, D. G. Park, N. Elmqvist, L. Kisselburgh, K. Ramani, skWiki: A multimedia sketching system for collaborative creativity, in: *Proceedings of the ACM Conference on Human Factors in Computing Systems*, ACM, Toronto, ON, Canada, 2014, pp. 1235–1244.