ActiveInk: (Th)Inking with Data

Hugo Romat, Nathalie Riche, Ken Hinckley, Bongshin Lee, Caroline Appert,
Emmanuel Pietriga, Christopher Collins

► To cite this version:

Hugo Romat, Nathalie Riche, Ken Hinckley, Bongshin Lee, Caroline Appert, et al.. ActiveInk: (Th)Inking with Data. CHI 2019 - The ACM CHI Conference on Human Factors in Computing Systems, May 2019, Glasgow, United Kingdom. 10.1145/3290605.3300272 . hal-02020272

HAL Id: hal-02020272

<https://hal.science/hal-02020272>

Submitted on 15 Feb 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

ActiveInk: (Th)Inking with Data

Hugo Romat
Microsoft Research
Redmond, WA, USA

Nathalie Henry Riche
Microsoft Research
Redmond, WA, USA

Ken Hinckley
Microsoft Research
Redmond, WA, USA

Bongshin Lee
Microsoft Research
Redmond, WA, USA

Caroline Appert
Univ. Paris-Sud, CNRS
Université Paris-Saclay
Orsay, France

Emmanuel Pietriga
INRIA
Université Paris-Saclay
Orsay, France

Christopher Collins
Microsoft Research
Redmond, WA, USA

ABSTRACT

During sensemaking, people annotate insights: underlining sentences in a document or circling regions on a map. They jot down their hypotheses: drawing correlation lines on scatterplots or creating personal legends to track patterns. We present ActiveInk, a system enabling people to seamlessly transition between exploring data and externalizing their thoughts using pen and touch. ActiveInk enables the natural use of pen for active reading behaviors, while supporting analytic actions by *activating* any of these ink strokes. Through a qualitative study with eight participants, we contribute observations of active reading behaviors during data exploration and design principles to support sensemaking.

CCS CONCEPTS

• **Human-centered computing** → **Touch screens; Visualization systems and tools**; • **Applied computing** → **Annotation**;

KEYWORDS

Digital pen, ink, visualization, active reading.

ACM Reference Format:

Hugo Romat, Nathalie Henry Riche, Ken Hinckley, Bongshin Lee, Caroline Appert, Emmanuel Pietriga, and Christopher Collins. 2019. ActiveInk: (Th)Inking with Data. In *CHI Conference on Human Factors in Computing Systems Proceedings (CHI 2019)*, May 4–9, 2019, Glasgow, Scotland UK. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3290605.3300272>

1 INTRODUCTION

Pen and touch technology with high-resolution displays, such as the Microsoft Surface Studio or Apple iPad Pro, are becoming mainstream. These devices offer much potential

for knowledge workers, affording large workspaces to handle multiple charts and documents, and natural interactions to manipulate and annotate them. Visualization researchers have started to explore this technology in systems tailored for pen and touch interactions (e.g., SketchInsight [15], VizDom [7], DataInk [38]). But these previous efforts have focused much less on how to enable people to externalize their thoughts as they engage in sensemaking activities.

Previous studies showed that a key affordance of analog pen and paper is to externalize one’s thoughts and capture fleeting ideas via annotating content and taking notes [18, 20, 24, 32]. Externalizing thoughts has many cognitive and social benefits [13] such as lowering working memory load, supporting idea reformulation, and providing common ground to share insights with others. These affordances are particularly important for sensemaking and data exploration tasks, as people need to keep track of numerous insights found while browsing through the data, which typically involves many visualizations and documents [23].

We contribute ActiveInk, a system that enables people to actively engage and think with data. It allows for a seamless transition between interacting with data and documents, and externalizing thoughts using a digital pen (Figure 1). We designed and evaluated two methods of working with ActiveInk. First, *prefix* provides pen modes and bimanual interaction to allow for switching between *ink pens* and *action pens*. Second, *postfix* provides ink for all pen strokes, and enables people to *activate* any ink stroke to perform analytic actions. For example, an annotation made at one point during the exploration process, such as underlining sentences in a document or circling a region in a map, may be activated at any time to filter, highlight, or color the underlying data. All actions propagate to all visualizations and documents in the workspace.

A study comparing these two approaches to a baseline of ink + touch (no actions) revealed benefits and drawbacks to each approach. We also observed that annotating data and taking ink notes were a common part of the analysis process across all conditions, but that ink actions, when available, supplanted some forms of annotation in active reading.

ACM, 2018. This is the author’s version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version will be published in *CHI 2019, May 4–9, 2019, Glasgow, Scotland UK*.

CHI 2019, May 4–9, 2019, Glasgow, Scotland UK

© 2019 Association for Computing Machinery.

<https://doi.org/10.1145/3290605.3300272>

Figure 1: ActiveInk affords smooth transition between using a digital pen for high-precision selections of data and for externalizing thinking via notes and annotations. Ink strokes are leveraged to perform operations on underlying data.

Informed by the study results, we expanded ActiveInk with a hybrid interface, providing pen modes as well as a *magic ink* pen, which provides ink by default, but ink strokes can be activated. We enhanced the analytic capabilities of ActiveInk to provide new functions such as activating a sketched correlation line to trigger a regression or hatching an interval on an axis to filter corresponding values. The resulting version of ActiveInk recognizes handwritten content as people take notes, allowing for words and shapes to be searched, filtered, and retrieved to facilitate revisitation.

In summary, our work contributes the following: (1) A novel set of pen-enabled actions for interacting with visualizations, documents, and images in a coordinated workspace; (2) A comparison of inking, *prefix* actions and *postfix* actions for exploratory analysis tasks; and (3) A hybrid design for pen + touch analysis which provides features specifically targeting the types of operation needed to support thinking during exploratory data analysis. These contributions bring people closer to their data, enabling them to interweave reading and acting on data. ActiveInk provides for seamless annotation, note-taking, and invocation of analytical operations, without even putting the pen down.

2 RELATED WORK

This research sits at the intersection of information visualization and human-computer interaction, building on previous designs of pen and touch interfaces for visualization, general

digital pen approaches for annotation and note-taking, and findings about the role of externalizing knowledge in solving data analysis tasks.

Pen and Touch Visualization Systems

Visualization researchers have recognized the potential of leveraging input modalities beyond mouse and keyboard in the past decade [14]. Advances in display technology and, in particular, the ability to offer simultaneous pen and touch on high-resolution displays have made possible new approaches to interact with visualizations.

There have been several approaches to design a coherent set of touch gestures for interacting with visualizations on tablets [8, 11, 25–27], but these do not fully take advantage of the different inking capabilities. Visualization systems such as SketchInsight [15], PanoramicData [40], and Viz-Dom [7] offer novel experiences leveraging pen and touch interactions for data analysis tasks. The research focus of these systems is on the design of interactions to create and interact with visualizations. Pen interactions often support two modes: free-form note-taking and predefined gestures to invoke commands. For example, sketching a particular shape (e.g., a circle) creates a specific type of visualization (e.g., a pie chart) in SketchInsight; drawing a line between two visualizations activates coordination through brushing and linking in PanoramicData. These interactions require learning and remembering a potentially large set of gestures.

Although free-form inking supported by these systems enables analysts to externalize their thoughts as they interact with the data, none of these ink marks are leveraged for further interaction.

A recent system, DataInk [38], enables people to sketch arbitrary shapes on a canvas and bind their visual properties (color, size, thickness) to data. While this research focused on creating expressive visualizations for storytelling and art purposes, it demonstrates one way to leverage the unique strengths of using a digital pen for visualizing data. We aim to provide more access to these affordances in ActiveInk.

Unique Affordances of Digital Pens

A recent study by Riche *et al.* [24] compiled and compared activities of people with an analog pen and using mainstream digital pen devices. They extracted a set of pen affordances, several of which are unique to digital pens. A key affordance of digital pens for working with visualizations is the accuracy for direct pointing and dragging, providing the ability to compose precise selections by sketching complex shapes on top of visualizations, while reducing the occlusion caused by touch. Digital pen systems have other capabilities beyond analog pens, most notably the ability for the same pen to perform different functions in quick succession.

Multiple affordances of an analog pen also appear to transfer to a digital pen [24] such as externalizing thoughts and capturing fleeting ideas. Research in human-computer interaction studied and leveraged these phenomena, in particular, for making annotations during active reading of documents [19, 20, 31], taking notes [10], gathering information [9, 35], generating ideas [17], and sharing ideas with others [41]. However, we are not aware of works which specifically focus on studying and supporting externalization using a pen for data analysis and sensemaking scenarios.

Externalization in Visualization

Externalizing thoughts through annotations on content, taking notes, and drawing diagrams has multiple benefits when engaging in long and complex tasks [13]. Externalizing thoughts enables people to limit their working memory load [29], articulate and reformulate thoughts which can lead to substantial improvements in understanding and retention [21], or share their thinking with peers to generate a shared object of thoughts to support communication and decision-making [5]. Indeed, Kidd hints that the act of annotation itself creates knowledge in the knowledge worker, which may be more important to the process than the marks that remain on the page [12].

These benefits are particularly relevant for visual analytics and sensemaking activities. The sensemaking process is a loop [23] involving phases of analysis and interpretation of visualizations and documents to gain insights, and phases

of revisitation to compare and contrast these insights to generate and investigate hypotheses. Externalizing thoughts is critical to facilitate revisitation and keep track of insights and hypotheses, especially over long periods of time.

Recent research started to study active reading of visualizations. Walny *et al.* [37] conducted a series of studies on active reading with both printed and digital visualizations. They demonstrated that active reading behaviors previously observed on textual documents (*e.g.*, underlining, writing notes in margin) occurred when people interpreted visualizations, and that such behaviors transferred in the digital world. While they did not consider complex sensemaking activities involving multiple visualizations and documents, their results suggest that inking with a pen on or around visualizations is an activity that people naturally do when working with visualizations (and given the ability to do so). This work builds upon their findings and reveals externalization strategies people employ during sensemaking activities using a digital pen on interactive data visualizations.

3 ACTIVEINK

Our goal with ActiveInk is to empower people to think with data. We aim at leveraging the digital pen to enable people to interweave *high-precision interactions* performed on visualizations and documents with *externalizing their thoughts* during the exploration process, as they would do with a regular pen on documents and blank paper. In the following, we define *note* as ink separated from data views, *annotations* as ink on data views and *actions* as ink that changes the visual appearance of data.

ActiveInk does not target expert data analysts as they would likely require extensive computations, analytic functions, and the ability to handle large amounts of data (better supported by systems such as Vizdom [7]). Rather, we seek to exemplify a more fluid sensemaking experience than what exists today for knowledge workers and data enthusiasts. We designed ActiveInk to primarily support thinking as people make sense of their data. We describe below the general principles that guided our design, and ActiveInk's key interface components.

Design Principles

We followed three driving principles to maximize opportunities for people to externalize thoughts while retaining the ability to interact with visualizations, minimizing the physical and cognitive costs of switching between the two.

D1. Support sensemaking through interaction with a set of heterogeneous information.

While many visualization systems focus on interacting with data visualizations, our goal is to support the broader activity of sensemaking, integrating documents and images as well. This activity involves extracting and cross-referencing

information from multiple sources. Thus, ActiveInk should provide people with space to think [1], *i.e.*, support viewing multiple of these assets at once and enable individual spatial organizations, where space can be given meaning.

D2. Use the digital pen for interacting with data and for externalizing thoughts.

A pen is more precise than a finger and offers a more natural interaction than a mouse for tasks such as drawing. Composing a selection in a visualization often relates to drawing as it involves the creation of complex arbitrary shape enclosing a set of points while avoiding others. Having the pen already in hand also maximizes the opportunities for externalizing thoughts (annotation, note-taking). Thus, we designed a set of pen actions to enable analysis. To reduce the cost of switching between analysis and externalization, we propose two interface strategies (Figure 2), which operate on an infinite canvas:

- (1) *prefix*: leveraging bi-manual interaction, in which people use their non-dominant hand to select an action for the pen they hold in their dominant hand; actions are directly performed as the pen is used;
- (2) *postfix*: the pen always lays ink on the screen; tapping ink strokes *a posteriori* with a finger reveals a menu of actions to perform on the data associated with that stroke. Activated strokes remain on canvas, allowing to recall and undo local actions at any time.

D3. Avoid requiring memorization of gestures.

Previous research on active reading and annotation has shown that the form of annotations varies between individuals, and even across tasks within individuals. A circle in one place may not serve the same annotative function as a circle in another, though similar annotations in spatio-temporal proximity often have the same intent [20]. Learning a collection of gestures may be difficult and recalling them may disrupt the flow of annotation. Instead, we aim to provide complete freedom of pen inking, and to support actions on the ink strokes that people naturally make.

Figure 2: Different sequences of interactions to perform an action such as paint on the data underlying the ink stroke.

Unified Infinite Canvas

ActiveInk is built around a zoomable infinite *canvas*, on which people can drop visualizations, documents, and images, which we collectively refer to as *views*, and rearrange them. The ability to lay out the workspace, reorganize and manipulate views was previously shown to be an important component of active reading of visualizations [37].

Freeform inking is provided anywhere on the canvas and on views alike. Ink on a view stays with the view when moved. Interaction with the canvas and views uses two-finger pinch to zoom. Single finger slide moves views or pans the canvas.

To provide powerful analysis across visualizations, the views imported into the ActiveInk canvas are tightly coupled. Each view responds to operations such as filter or highlight performed on the others. To maximize usability and minimize the cognitive effort of tracking coordination, we chose to coordinate all views of the data on the canvas. The data is loaded from tabular data file formats such as CSV, and is stored in-memory as a JSON data structure. In this first prototype we do not use a visualization toolkit, but rather rely on a set of basic, common chart types implemented from scratch in our environment for the sake of easy customization. ActiveInk currently supports binned histograms, scatterplots, maps, and text documents. We selected this subset for the diversity of data, interactions, and insights they can offer. However, adding additional types of visualizations such as line charts or density plots is straightforward.

Pen Actions

ActiveInk supports the set of operations on views illustrated in Figure 3. The following scenario, illustrated in Figure 4, describes the pen actions of ActiveInk and how one can use these operations for sensemaking using the *prefix* strategy.

Emma is a data journalist writing a piece about the best states in the U.S. to start a family. She has collected data about each state, such as school quality, air pollution, and median home price, as well as documents about this topic. She organizes the most interesting views in her canvas, moving them with a finger and pinching them to resize. She annotates salient insights with *ink* as she encounters them. For example, she notes that the school quality index mostly correlates with math scores in Grade 8. She uses *cut* to slice the binned histogram of house prices, as the distribution is skewed. Providing a healthy environment being critical, she decides to remove all states that have higher air pollution by switching to the *remove* pen and using it on dark regions of the map. Data is removed from the entire canvas, and the scales in scatterplots and histograms are updated in response.

Figure 3: Visual effects produced by our 8 different types of pens in *prefix* (a). In *postfix* (b), tapping an ink stroke reveals the activation menu. Sliding a finger from the center to an action previews its effect, while lifting the finger triggers it.

Since school quality is an important factor to raise a child, Emma opts to *hide* the lower quadrant of the scatterplot, deemphasizing those states in all views. She then chooses a blue color to *highlight* the higher portion of happiness score. She notices that one of them, Vermont, pops up in the document and thus reads the corresponding paragraph. To identify all data about Vermont, she switches to a purple color to *paint* and *label* it from the text. Since Emma wants to have more details on house prices in Vermont, she uses the *copy* instrument and lassos the interval containing Vermont. She sees Vermont is in the lower range, concluding that it is the place to be!

Prefix Interface: Switching between Types of Pen

The *prefix* interface offers the pen actions described above as different types of instruments people can acquire with the pen (Figure 3). This paradigm of different *modes of interaction* is usually found in mainstream interfaces today, although modes generally appear in a menu or ribbon positioned at the top. In contrast, ActiveInk presents it on the non-dominant hand side of the workspace to afford bi-manual Thumb+Pen interaction [22], in which a person taps on modes with the thumb of one hand while using the pen in the other hand. This design may lower the cost of interaction for mode switching [3].

The ActiveInk *prefix* interface uses a paradigm familiar to most people and provides a comprehensive view of all actions at all times, which is important for non-experts [30]. This paradigm is also efficient for accomplishing actions in batches. However, people need to decide on mode before performing actions, and interweaving different types of activities may induce numerous mode switches. Forgetting to switch modes may result in mistakes to undo later (e.g., people want to rapidly jot down thoughts next to highlighted points in a scatterplot but the pen highlights more points instead). Action pens in *prefix* lay down dashed ink to mitigate mode errors, and the interface includes buttons for global

undo and redo to correct mistakes. The *clear* pen $\times$ is used to clear all visual properties of data elements resulting from prior interactions (highlight, paint, etc.), affording some level of local undo.

Postfix Interface: Activating Ink Strokes

The *postfix* interface offers the pen actions described above through an in-place menu invoked by tapping an ink stroke (Figure 3, right). Activated ink strokes remain on canvas and can be deactivated to undo the action. We hypothesized that this design would stimulate externalizations as the pen always lays down ink first without requiring to switch back to this mode after interacting with the data.

We organized the actions in the menu along three semantic axes, inspired by the design of the ZoomCatcher menu [39].

- 1 **Visual saliency** of data elements that one can increase with *Highlight* and *Paint*, and decrease with *Hide* and ultimately *Remove*.
- 2 **Display data** about the underlying elements with *Label* displaying the name of data entities and their values in histograms.
- 3 **View manipulations** that do neither affect the data nor propagate to other views, including *cut* and *copy*.

Once the menu is invoked by tapping a stroke, one can preview the effects of any action before executing it by sliding their finger from the central ink action to any other menu item. Lifting the finger triggers the action and fades out the ink stroke to limit visual clutter. Note that actions that are unavailable (e.g., due to no underlying data to highlight) are greyed out. Erasing an activated (faded-out) stroke undoes the action and reverts the stroke to a regular ink (that can potentially be activated again, or fully erased). This mechanism allows for local undo of actions during the exploration.

The *postfix* interface uses a paradigm that is less familiar to people, and certainly has a steeper learning curve, as all actions are not visible at all times. Given the additional

Figure 4: Depiction of the canvas of a data journalist organizing their workspace and annotating findings (1-ink , 2-cut , 3-remove) and interacting on the data to identify salient subsets (4-hide , 5-highlight , 6-paint , and label , and 7-copy) a subset of a histogram).

cost of tapping after each ink stroke, this paradigm may slow down actions performed in batches. However, it enables the smooth interweaving of interactions with data and ink externalizations without mode switch (e.g., paint a set of points and hand write the reason for this action next to them). Additionally, actions can be carried out at any time: an underline, arrow, or circle made early during an initial read of the data can later be activated by tapping, thus turning externalizations into potential interaction locations. Since activated strokes remain on the canvas, we also hypothesized that they would provide additional benefits for analytic provenance — assist people in recalling their process to gain insights on the data. Finally, the *postfix* interface has fewer tool icons, and thus is less cluttered.

Implementation

ActiveInk is built as a web-based application in JavaScript and runs on a NodeJS server. The vector graphics canvas is implemented using Paper.js [16] and the visualizations are created with d3.js [4] and BubbleSets.js [6]. To best associate ink strokes with data, ActiveInk tests for data item enclosure or intersection. This allows for marks such as circles and arrows to appropriately associate with the intended data.

4 QUALITATIVE STUDY

To gain insights on the externalization process during sense-making using a digital pen, we conducted a qualitative observation study. We designed the procedure to investigate the following research questions:

Q1. How do externalizations manifest during sensemaking with digital inking?

Our first question is whether results from Walny *et al.* [37] will transfer to more complex tasks involving extracting and comparing insights from multiple visualizations and documents to make sense of the data.

Q2. How do externalizations manifest during sensemaking with ActiveInk?

We seek to observe if and how annotation and note-taking behaviors change when offered the possibility to interact with the data to visually encode it (e.g., highlight). In particular, a key question is whether the cost of switching mode between interaction with data (e.g., highlight, filter) and inking in *prefix* may discourage spontaneous externalization of thinking. In contrast, while *postfix* may encourage such externalization, it uses a less familiar interaction paradigm where actions are not visible at all time, which may prove more difficult to use.

Study material and screenshots of canvases created by participants are available as supplemental material.

Study Design

We used a within-participants design as prior work on active-reading [18, 20, 37] indicated that individual variability would be high given the diverse range of strategies for sensemaking and behaviors for externalization. Thus, participants experienced the following three interfaces in a single session.

Ink. The condition explored sensemaking using only digital ink and touch. Participants could browse through a set of views (visualizations and documents), adding and organizing them in the workspace with touch. To inspect underlying data, participants could tap and hold with a finger on a visual element to display its label and data values in all views. Inking was available anywhere on the canvas, including on views. This condition served as our baseline, enabling us to build knowledge on active reading and externalization behaviors without any analytic functions linked to pen interactions.

ActiveInk Prefix and Postfix. In the second and third conditions, participants used ActiveInk with *prefix* and *postfix* interfaces, as described above. Inking and touch capabilities were also provided as in the ink condition. As we hypothesized an order effect (due to fatigue or learning), we counterbalanced the presentation order of these two conditions. Participants also had to complete different scenarios in each of these two conditions. Assignment between a scenario and a condition was counterbalanced across participants in order to avoid observing an effect of the dataset instead of observing an effect of the ink condition.

Participants

We recruited 8 participants (2 females, 6 males; 7 in 30s and 1 in 20s) from a large software company (over 60000 employees), screening them to cover different age ranges, genders, and roles (team managers, administrative assistants, developers, data scientists, researchers). We screened participants for at least one year of experience with a pen-enabled device and recent experience reading and creating simple charts from data. Only P5 reported data analysis is not part of his/her job. We piloted the study with 2 people to streamline the training and fix minor usability issues.

Method

We used a Microsoft Surface Studio, a 28" screen with a resolution of 4500×3000 pixels, supporting multitouch and pen input. Participants were encouraged to adjust the orientation and incline of the device for comfortable pen and touch interactions. The study was held in a quiet room; experimenters observed live, from behind a space divider, using a camera.

Data and Task. As our goal was to observe externalization during sensemaking, we designed high-level tasks involving

browsing through a set of views, making hypotheses, and investigating queries to understand the data. We prepared 3 similar analytic scenarios based on identifying a subset of states in the U.S. In the ink condition, participants searched for states where tobacco and alcohol use was correlated to accidents and overdoses. In the *prefix* and *postfix* conditions we counterbalanced across two scenarios: identifying the best and worst states to retire, or start a new job and raise a family. A different multivariate dataset was provided for each scenario, containing factors such as nursing home quality (retirement) and home prices (new job). Examples appear in Figures 1 and 4. Participants were instructed to imagine they were conducting research to write an article about their findings, and should create a canvas that would enable them to recall these findings in several weeks. To limit the duration of the session, we selected a subset of visualizations and documents (9 histograms, 9 scatterplots, 3 documents, and 3 maps) for each exploration task.

Procedure. An experimenter first instructed participants to sign a consent form and fill out a demographics questionnaire. The study was structured in three phases, one per condition. For each condition, the experimenter briefly described the interface and dataset at the start of the training. To cover the interaction mechanics, the experimenter instructed participants to perform each action available in the interface (pen and touch), explaining the interaction and resulting effects, and informing participants of efficient interaction techniques if needed (e.g., bi-manual interactions). This training lasted about 10 minutes. Since ActiveInk features a large number of actions, these conditions included an additional 5 minute practice task, in which participants completed a brief analysis task while asking questions and receiving experimenter guidance as needed.

After training, participants completed the main data exploration task for a maximum of 15 minutes. After each task, participants completed a questionnaire and had a brief interview with the experimenter about their likes, dislikes, and wishes. Participants were instructed to take a 5 minute break after the second condition. After experiencing all conditions, participants answered a preference questionnaire and verbally explained the advantages and drawbacks of each technique. The entire experiment lasted two hours. Participants received a \$150 gift card as compensation.

Data Collection. We logged pen interactions, and used a screen video capture and video-audio recordings. At least two experimenters were present for each session: one person interacted with the participant while the other took notes.

Figure 5: Inking and action sequences per participant across conditions. Each block reflects a 15 minute analysis task. Participants carried out the most annotations in the ink condition, followed by *postfix*. The most common operations were highlight, paint, and label, which provide details on the data. Split bars in *postfix* indicate ink strokes used for multiple actions. Note that P8-*postfix* logs were not saved due to a technical error.

Results

Figure 5 depicts the interactions of each participant over the course of each exploration task. All canvases created in the study are in supplemental material.

Ink. In the first condition, six participants used inking extensively for making sense of data (Figure 5). The largest portion of ink across participants are *notes* separated from data (78.0% of strokes), as opposed to direct *annotations* on views (22.0%). Participants exhibited different behaviors: P1 and P5 laid down most ink, however, P1 made both notes and annotations while P5 only took notes. While P4 and P8 did not ink as much as others, their ink mostly consisted of annotations.

Figure 6 depicts externalizations in the ink-only condition. It is interesting to note that two participants (P1, P4 in Figure 6) used inking to coordinate multiple subsets of data permanently on the canvas. These behaviors correspond to the paint, highlight, and label actions of ActiveInk. We also observed the depiction of hypotheses about the data (e.g., drawing a correlation line) and externalized queries (e.g., drawing a threshold). Both of these annotations are illustrated in Figure 6-P1.

Most of the notes on the canvas were phrases summarizing findings, hypotheses, or describing the sensemaking process itself. However, we also saw instances of pictorial representations such as arrows (6 participants), or lines on the canvas to divide the workspace into regions (P3, P5). We also observed participants visually link a pictorial annotation in a visualization to a separate handwritten note. For example, P1 drew a box in red around the word “threshold” referencing the line sketched with the same pen in both scatterplots (Figure 6-P1). Other participants (P1, P2, P6, P7) used different colors, visually connecting notes to relevant annotations.

Figure 6: Examples of participants' externalizations during sensemaking with ink. P1 used color to connect notes and annotations, while P4 drew lines to connect data items across views.

Externalizations with ActiveInk. In contrast to the ink-only condition, the amount of regular ink in both ActiveInk conditions decreased, with fewer annotations. The prevalence of notes did not decrease as much: 8% lower in *prefix* and slightly higher in *postfix*. The total number of annotation + action strokes was consistently between 19–26% across all conditions. This may indicate that actions serve the role of some annotations in ActiveInk.

While there were fewer annotations than in the ink condition, we noted several interesting types of annotations created after operating on the data: (1) depicting how views

Figure 7: Examples of participants’ externalizations during sensemaking with ActiveInk, including batch highlighting with *prefix* on the left, and interwoven actions and note-taking with colors using *postfix* on the right.

relate after applying cut or copy (e.g., copying a subset of a view), (2) depicting insights (e.g., correlation between two data dimensions), (3) identifying data of interest (e.g., circling), and (4) characterizing data of interest (e.g., good vs. bad candidates). In the ink condition, we saw annotations to connect specific data points across multiple views. We did not observe this type of annotation with ActiveInk since actions such as highlight and paint fulfilled this role.

The ink stroke log analysis (Table 1) indicates that *postfix* encourages more annotation than *prefix*. In observing the use of *prefix*, we saw that participants chose a pen and repeatedly made short strokes on items one by one, whereas with *postfix* they more commonly used enclosures to select groups of items, then activated the ink.

The most frequent action performed across both ActiveInk conditions was highlight (187), followed by label (55), paint (50), filter (30), cut (6), copy (3), hide (3), and remove (2). Video analysis shows that cut and remove were primarily used on skewed distributions and outliers in scatterplots and histograms. In the *prefix* condition, the clear function was used to remove actions 47 times, and undo was used 39 times. Undo and clear were not provided in *postfix* as ink actions were locally reversible at all times. All participants commented that ActiveInk analytic functions on data were useful in contrast to the ink condition.

While a higher fraction of strokes were used for note-taking in *postfix*, the note-taking behavior was consistent in

Stroke Type	Ink	<i>prefix</i>	<i>postfix</i>
Notes	1082	78.0%	983 70.1%
Annotations	305	22.0%	59 4.2%
Actions	0	314 22.4%	89 9.1%
Errors	0	46 3.3%	6 0.6%
Total	1387	100%	1402 100%

Table 1: Ink stroke analysis. Errors are data action on empty canvas.

both ActiveInk conditions. For example, P7 wrote hypotheses on top of the canvas, interacted with data visualizations to investigate them, and wrote answers next to original notes in both conditions. Similar to the ink condition, participants created colored legends (e.g., Figure 7-P2) in both ActiveInk conditions. Overall, we observed meaning assigned to color in 75% of all exploration tasks across participants and conditions. We observed that participants generally followed one or more of these strategies: 1) write the data dimension name or value, 2) use the same color to mark the data subset, and 3) use a pictorial representation (e.g., underline, check mark) to connect notes in one place to data in different visualizations.

Sensemaking strategies with ActiveInk. We observed two different strategies that reflect advantages of each interface (Figure 7). P2 (Figure 7-P2) seamlessly interweaved paint action on data item, and in-place annotation with the same color to create a pictorial legend about this data and value.

Figure 8: Advanced actions added to ActiveInk, informed by results of our qualitative study.

He sketched glyphs such as check marks and thumbs up to indicate positive aspects, and sad faces and thumbs down to indicate negative aspects. This strategy worked well with *postfix* as the pen “automatically” switched to ink after performing an action on the data. In contrast, P6’s strategy was to do actions in batches (Figure 7-P6). Using the ink pen, P6 first handwrote a legend to associate certain data values to positives and negatives, then switched to the highlight pen to perform highlights in batches, only switching colors. This strategy worked well with *prefix*, as it did not require tapping strokes, but rather pick a single type of pen to apply successive strokes on the views.

We observed differences between participants but each followed a personal sensemaking strategy, generally remaining consistent across conditions (Figure 5). For example, P1 used inking extensively in the first condition and continued to take more notes than others in *prefix* and *postfix* while most of the annotations were replaced by pen actions. In contrast, P4 mostly relied on annotations in ink condition, which were replaced by pen actions (mostly highlight) in ActiveInk.

Most participants (6/8) attempted to lay out visualizations and documents on the canvas so they were all in view without requiring zooming or panning. P7 made use of the larger surface, storing handwritten notes on hypotheses and results out of screen. Two participants used a lower corner to place temporary visualizations. P8 used the histograms as one would use dynamic queries in conventional coordinated view interfaces, as a way to filter the data according to different attributes. P8 is also the only one who leveraged multiple tabs to segment the workspace by different questions investigated (e.g., one canvas for identifying the best states where to retire, another for the worst ones).

Ratings and Preferences. Participants praised the quality of both interfaces. For example, P1 said “I really like it, it’s the optimal for this type of task really.” (referring to the general principle of ActiveInk, independently of *prefix* or *postfix*).

Note that while our questionnaire did not have an option to rate as equal, three participants verbally indicated that they did not have a clear favorite. Overall, six participants preferred *prefix* and two preferred *postfix*. In *prefix*, six participants used bimanual interaction, five of whom preferred *prefix* over *postfix* overall.

Results of the subjective experience questionnaire administered after each condition revealed similar ratings for all conditions, although *postfix* rated slightly higher than *prefix* on most measures. Most (6/8 in *prefix* and 8/8 in *postfix*) felt efficient with ActiveInk. Participants made more notes and annotations with *postfix*, while with *prefix*, they carried out more actions. In *prefix*, at least 3.3% were mode errors (actions on the canvas) (Table 1). Although, note that we could not detect mode errors resulting from accidental inking. Four participants specifically mentioned mode errors as frustrating, e.g., P3: “I keep just forgetting which one I’m in,” P7: “When I was highlighting and wanted to go back to my notes I forgot to click here.”

In *postfix*, the menu was activated without previewing or applying any action 6 times, which we also counted as errors. We also observed three participants experiencing usability issues with the *postfix* menu. In particular, P1 said: “I would like something faster, like tap and choose rather than tap, wait for animation, and drag to choose.” While dragging on the menu was intended to preview action results before executing them, these previews were not frequently used (only six instances from three participants). Several requested to see all available actions at once. This may be especially important during the learning period as participants are not familiar with the semantics of the actions available and thus do not fully comprehend their organization in the menu.

Wishes. Three participants mentioned the potential of a hybrid interface, integrating the best of both approaches. P3 said, “Can you make something where some people can choose the tool first and other people can select things and change them?” while P9 suggested using physical buttons on the

pen to activate the most common actions. P5 suggested the ability to see the underlying data in a table, and to be able to change data values directly in the interface. P8 commented on the need for transient pen highlights “*I would like to drag my ink across the thing and ‘boom’ things light up.*” In addition to local undo provided by erasing an activated ink in *postfix*, P2 and P6 suggested a global undo button would also be helpful. P2 and P6 also suggested functions to align or snap views together for neater canvases. Finally, both P1 and P5 suggested sketching dividing lines on the canvas to create different sets of coordinated views (in addition to different tabs).

Summary and Discussion

Our results suggest that the results of Walny et al. on active reading of one visualization [37] transfer to sensemaking tasks involving multiple visualizations and documents. When ink actions are available, people favor visually marking the data using actions rather than annotating with regular ink, though note-taking was present in all conditions. We observed that several annotations were used to add information not built into the action tools (e.g., sketched glyphs, semantic use of color). Annotations were reused by all participants in describing their findings and strategies to the experimenter, though we did not specifically test for the benefit of inking for recall. While we believe there are sensemaking benefits resulting from the low cost of externalizing thoughts with digital inking, more investigation is needed to check whether this behavior also happens in conventional systems using mouse and keyboards.

Most participants used notes, physically separated from the data, to record their questions and hypotheses, and evidence to confirm or reject them. Participants linked their handwritten notes to data by either rewriting the data dimension name or value, using the same ink color as used for data highlights, or using a pictorial representation (e.g., a check mark, an arrow). These techniques reflect the creation of ad-hoc visual indices to connect and cross-reference related information from different sources. This manual and iterative process may explain why most participants attempted to keep everything in one view, and suggests the need to better support management of large canvases or multiple pages of information.

Implications for Design

As evidenced by the ink stroke analysis, the *prefix* and *postfix* interfaces lead to different behaviors: *prefix* supported rapid sequential actions, while *postfix* better integrated with annotation and note-taking. This implies a better solution would be to support both paradigms. Thus, we created a hybrid version which offers the same set of pens available in the *prefix* interface, as well as the ability to activate strokes laid

down with the ink pen and marker using the *postfix* activation menu. As before, activated ink strokes remain on the canvas, constituting a visual trace of the action performed, and a means to undo the action later on.

The observations we made about note-taking during sensemaking with ActiveInk prompted us to add a 4th design principle:

D4. Make externalizations useful.

We observed that, as people explore data, they took notes separated from the data visualizations or documents to capture their hypotheses or summarize their findings. These ink strokes contained explicit references to data dimensions or values through textual content, color, or associated shape to enable cross-referencing insights in multiple sources (notes, data visualizations, and documents). ActiveInk should leverage these externalizations for further analysis.

Advanced Actions to Support Sensemaking

The result of our qualitative study informed a set of 6 new functions depicted in Figure 8, inspired by our observations and comments from our participants.

Compute statistics: Sketching hypotheses about statistical properties of the data directly on the visualization is common (e.g., sketching a correlation line in a scatterplot). While laying ink is a mean to quickly mark potentially important insights discovered serendipitously while exploring the data, getting back to these and assessing their correctness is important when drawing conclusions about the data. For this reason, we added the ability to compute such statistics by tapping on a stroke and selecting the statistics action (i.e., adding a new options in the *postfix* menu). Note that ActiveInk currently supports computing a correlation line in scatterplots and averaging the values of a set of data items in histograms as we mostly observed these behaviors in the study.

Interactive legends: We also observed participants attempting to draw on the map legend (and sometimes tap it with a finger). We incorporated the ability to sketch on the legend to perform actions (highlight, paint, label, hide, etc.) on the corresponding data elements.

Ink lenses: Study participants surfaced the need for transient and interactive pen interactions akin to brushing in conventional data visualization systems. We added the ability to create and activate any set of strokes on empty canvas with functions such as highlight or paint. Moving these activated strokes with a finger over views immediately performs the corresponding action on the underlying data, effectively acting as lenses on the data [36] and acting as dynamic queries [34] when activated with filtering.

Recognizing ink content: Since we saw many references to data dimensions and values in the notes participants took, we added handwriting recognition to ActiveInk. This capability

turns handwritten notes into content coordinated with the views. In particular, ActiveInk now offers the ability to tap a written word and perform an action on related data. This capability is inspired from the behavior we saw participants exhibit when referring to data dimensions or values in their notes to connect them to the data. For example, one would observe states with low unemployment, identify Texas and California. Highlighting Texas to further investigate its values along other dimensions, they would make a note to get back to California later. ActiveInk facilitates this process as one would simply tap the written word California, select paint from the *postfix* menu and see the corresponding data items in visualizations and text using the same ink color. Note that handwriting recognition is performed with the MS Ink Analysis API [33].

Search similar ink shapes: Since we observed people draw the same glyph on different views, we added the ability to search for similar sets of strokes making up a glyph. When a search is triggered on a glyph, ActiveInk retrieves similar sets of strokes and brings them, and their underlying content, close to the location of the search action. This facilitates cross-referencing and may prove useful for a large canvas. Strokes are associated with glyphs using a spatiotemporal distance function and threshold, while glyphs are associated with each other using the \$N\$-recognizer [2].

Show and edit underlying data: Participants requested the ability to see the underlying data table for a subset of data, and one participant suggested she would like to edit the data directly. We added the *show data* action to reveal the underlying data table for items associated with a pen stroke. When writing numbers in the table, the written values are recognized and the corresponding data items are updated across all views.

5 CONCLUSION AND FUTURE WORK

ActiveInk represents a foray into active reading applications for heterogeneous data including visualizations, maps, and documents. It blends the operations of externalizing thoughts with analytic actions in a seamless combination of functions to be carried out with digital pen+touch. The hybrid interface of ActiveInk integrates two ways of working with pen actions, which our study revealed to be complementary for sensemaking activities. *Prefix* inking is useful for batch operations, while *postfix* ink activation is useful for interleaving multiple actions with externalizing thoughts.

This work offers several opportunities for extension. The ink recognition system we used does not distinguish between classes of marks, such as enclosures, underlines, and arrows. Our data selection model associates marks with any data items enclosed or intersected by an ink stroke. Distinguishing types of mark could be interesting to implement a finer ink-data association strategy.

ActiveInk was made to demonstrate the possibilities of thinking with ink, but does not scale to a large number of views or with very large volumes of data. Addressing scalability issues requires further development, and implementing the ability to save, recall, and share analysis canvases. It also implies adding support for a wider diversity of charts, which could be achieved by moving to, e.g., Vega-lite [28].

To better understand the impact of ActiveInk on sensemaking and externalization, including any potential effects on the depth and quality of insights, we would like to run a longitudinal study in which participants analyze their own data using the hybrid interface. We suspect performance may improve with long-term use, and sensemaking strategies may change as people feel more engaged with their data. Furthermore, it would be interesting to study the use of ActiveInk in collaborative synchronous and asynchronous sensemaking scenarios.

REFERENCES

- [1] Christopher Andrews, Alex Endert, and Chris North. 2010. Space to Think: Large High-resolution Displays for Sensemaking. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*. ACM, 55–64. <https://doi.org/10.1145/1753326.1753336>
- [2] Lisa Anthony and Jacob O. Wobbrock. 2012. \$N\$-protractor: A Fast and Accurate Multistroke Recognizer. In *Proc. Graphics Interface*. Canadian Information Processing Society, 117–120.
- [3] Caroline Appert, Michel Beaudouin-Lafon, and Wendy E. Mackay. 2005. Context matters: Evaluating Interaction Techniques with the CIS Model. In *People and Computers XVIII —Design for Life*, Sally Fincher, Panos Markopoulos, David Moore, and Roy Ruddle (Eds.). Springer, London, 279–295. https://doi.org/10.1007/1-84628-062-1_18
- [4] Michael Bostock, Vadim Ogievetsky, and Jeffrey Heer. 2011. D3 Data-Driven Documents. *IEEE Trans. on Visualization and Computer Graphics* 17, 12 (Dec. 2011), 2301–2309. <https://doi.org/10.1109/TVCG.2011.185>
- [5] Mauro Cherubini, Gina Venolia, Rob DeLine, and Andrew J. Ko. 2007. Let’s Go to the Whiteboard: How and Why Software Developers Use Drawings. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*. ACM, 557–566. <https://doi.org/10.1145/1240624.1240714>
- [6] Christopher Collins, Gerald Penn, and Sheelagh Carpendale. 2009. Bubble Sets: Revealing Set Relations with Isocontours over Existing Visualizations. *IEEE Trans. on Visualization and Computer Graphics* 15, 6 (2009), 1009–1016. <https://doi.org/10.1109/TVCG.2009.122>
- [7] Andrew Crotty, Alex Galakatos, Emanuel Zraggen, Carsten Binnig, and Tim Kraska. 2015. Vizdom: Interactive analytics through pen and touch. *Proc. of the VLDB Endowment* 8, 12 (2015), 2024–2027. <https://doi.org/10.14778/2824032.2824127>
- [8] Steven M. Drucker, Danyel Fisher, Ramik Sadana, Jessica Herron, and m.c. schraefel. 2013. TouchViz: A Case Study Comparing Two Interfaces for Data Analytics on Tablets. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*. ACM, 2301–2310. <https://doi.org/10.1145/2470654.2481318>
- [9] Ken Hinckley, Xiaojun Bi, Michel Pahud, and Bill Buxton. 2012. Informal information gathering techniques for active reading. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*. ACM, 1893–1896. <https://doi.org/10.1145/2207676.2208327>
- [10] Ken Hinckley, Ken Hinckley, Shengdong Zhao, Raman Sarin, Patrick Baudisch, Edward Cutrell, Michael Shilman, and Desney Tan. 2007. InkSeine: In Situ Search for Active Note Taking. In *Proc. of the SIGCHI*

- Conf. on Human Factors in Computing Systems*. ACM, 251–260. <https://doi.org/10.1145/1240624.1240666>
- [11] Jaemin Jo, Sehi L’Yi, Bongshin Lee, and Jinwook Seo. 2017. Touch-Pivot: Blending WIMP & Post-WIMP Interfaces for Data Exploration on Tablet Devices. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*. ACM, 2660–2671. <https://doi.org/10.1145/3025453.3025752>
 - [12] Alison Kidd. 1994. The Marks Are on the Knowledge Worker. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*. ACM, 186–191. <https://doi.org/10.1145/191666.191740>
 - [13] David Kirsh. 2010. Thinking with external representations. *AI & Society* 25, 4 (2010), 441–454. <https://doi.org/10.1007/s00146-010-0272-8>
 - [14] Bongshin Lee, Petra Isenberg, Nathalie Henry Riche, and Sheelagh Carpendale. 2012. Beyond mouse and keyboard: Expanding design considerations for information visualization interactions. *IEEE Trans. on Visualization and Computer Graphics* 18, 12 (2012), 2689–2698. <https://doi.org/10.1109/TVCG.2012.204>
 - [15] Bongshin Lee, Greg Smith, Nathalie Henry Riche, Amy Karlson, and Sheelagh Carpendale. 2015. SketchInsight: Natural data exploration on interactive whiteboards leveraging pen and touch interaction. In *Proc. of the IEEE Pacific Visualization Symposium*. IEEE, 199–206. <https://doi.org/10.1109/PACIFICVIS.2015.7156378>
 - [16] Jürg Lehm and Jonathan Puckey. 2018. Paper.js. <http://paperjs.org>
 - [17] Guang Li, Xiang Cao, Sergio Paolantonio, and Feng Tian. 2012. SketchComm: a tool to support rich and flexible asynchronous communication of early design ideas. In *Proc. of the ACM Conf. on Computer Supported Cooperative Work*. ACM, 359–368. <https://doi.org/10.1145/2145204.2145261>
 - [18] Catherine C. Marshall. 1997. Annotation: From Paper Books to the Digital Library. In *Proc. of the ACM Int. Conf. on Digital libraries*. ACM Press, 131–140. <https://doi.org/10.1145/263690.263806>
 - [19] Fabrice Matulic and Moira C Norrie. 2012. Supporting active reading on pen and touch-operated tabletops. In *Proc. of the Int. Working Conf. on Advanced Visual Interfaces*. ACM, 612–619. <https://doi.org/10.1145/2254556.2254669>
 - [20] Hrim Mehta, Adam James Bradley, Mark Hancock, and Christopher Collins. 2017. Metatation: Annotation as Implicit Interaction to Bridge Close and Distant Reading. *ACM Trans. on Computer-Human Interaction (TOCHI)*, Article 35 (2017), 41 pages. <https://doi.org/10.1145/3131609>
 - [21] Pam A Mueller and Daniel M Oppenheimer. 2014. The pen is mightier than the keyboard: Advantages of longhand over laptop note taking. *Psychological Science* 25, 6 (2014), 1159–1168. <https://doi.org/10.1177/0956797614524581>
 - [22] Ken Pfeuffer, Ken Hinckley, Michel Pahud, and Bill Buxton. 2017. Thumb + Pen Interaction on Tablets. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*. ACM, 3254–3266. <https://doi.org/10.1145/3025453.3025567>
 - [23] Peter Pirolli and Stuart Card. 2005. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. (2005), 2–4. https://analysis.mitre.org/proceedings/Final_Papers_Files/206_Camera_Ready_Paper.pdf
 - [24] Yann Riche, Nathalie Henry Riche, Ken Hinckley, Sheri Panabaker, Sarah Fuelling, and Sarah Williams. 2017. As We May Ink?: Learning from Everyday Analog Pen Use to Improve Digital Ink Experiences. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*. ACM, 3241–3253. <https://doi.org/10.1145/3025453.3025716>
 - [25] Jeffrey M. Rzeszutarski and Aniket Kittur. 2014. Kinetica: Naturalistic Multi-touch Data Visualization. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*. ACM, 897–906. <https://doi.org/10.1145/2556288.2557231>
 - [26] Ramik Sadana and John Stasko. 2014. Designing and implementing an interactive scatterplot visualization for a tablet computer. In *Proc. of the Int. Working Conf. on Advanced Visual Interfaces*. ACM, 265–272. <https://doi.org/10.1145/2598153.2598163>
 - [27] Ramik Sadana and John Stasko. 2016. Designing Multiple Coordinated Visualizations for Tablets. *Computer Graphics Forum* 35, 3 (June 2016), 261–270. <https://doi.org/10.1111/cgf.12902>
 - [28] Arvind Satyanarayan, Dominik Moritz, Kanit Wongsuphasawat, and Jeffrey Heer. 2017. Vega-lite: A grammar of interactive graphics. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2017), 341–350.
 - [29] Michael Scaife and Yvonne Rogers. 1996. External Cognition: How do Graphical Representations Work? *Int. J. Hum.-Comput. Stud.* 45 (Aug. 1996), 185–213. <https://doi.org/10.1006/ijhc.1996.0048>
 - [30] Joey Scarr, Andy Cockburn, Carl Gutwin, and Andrea Bunt. 2012. Improving Command Selection with CommandMaps. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*. ACM, 257–266. <https://doi.org/10.1145/2207676.2207713>
 - [31] Bill N Schilit, Gene Golovchinsky, and Morgan N Price. 1998. Beyond paper: Supporting active reading with free form digital ink annotations. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*. ACM, 249–256. <https://doi.org/10.1145/274644.274680>
 - [32] Abigail J. Sellen and Richard H.R. Harper. 2003. *The Myth of the Paperless Office*. MIT Press, Cambridge, MA, USA.
 - [33] Microsoft Cognitive Services. 2018. Ink Analysis API. <https://labs.cognitive.microsoft.com/en-us/project-ink-analysis>
 - [34] Ben Shneiderman. 1994. Dynamic Queries for Visual Information Seeking. *IEEE Softw.* 11, 6 (Nov. 1994), 70–77. <https://doi.org/10.1109/52.329404>
 - [35] Craig S. Tashman and W. Keith Edwards. 2011. LiquidText: A Flexible, Multitouch Environment to Support Active Reading. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*. ACM, 3285–3294. <https://doi.org/10.1145/1978942.1979430>
 - [36] Christian Tominski, Stefan Gladisch, Ulrike Kister, Raimund Dachselt, and Heidrun Schumann. 2014. A Survey on Interactive Lenses in Visualization. In *EuroVis - STARs*, R. Borgo, R. Maciejewski, and I. Viola (Eds.). The Eurographics Association. <https://doi.org/10.2312/eurovisstar.20141172>
 - [37] Jagoda Walny, Samuel Huron, Charles Perin, Tiffany Wun, Richard Pusch, and Sheelagh Carpendale. 2018. Active Reading of Visualizations. *IEEE Trans. on Visualization and Computer Graphics* 24, 1 (2018), 770–780. <https://doi.org/10.1109/TVCG.2017.2745958>
 - [38] Haijun Xia, Nathalie Henry Riche, Fanny Chevalier, Bruno De Araujo, and Daniel Wigdor. 2018. DataInk: Direct and Creative Data-Oriented Drawing. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*. <https://doi.org/10.1145/3173574.3173797>
 - [39] Haijun Xia, Ken Hinckley, Michel Pahud, Xiao Tu, and Bill Buxton. 2017. WritLarge: Ink Unleashed by Unified Scope, Action, & Zoom. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*. ACM, 3227–3240. <https://doi.org/10.1145/3025453.3025664>
 - [40] Emanuel Zraggen, Robert Zeleznik, and Steven M Drucker. 2014. PanoramicData: Data analysis through pen & touch. *IEEE Trans. on Visualization and Computer Graphics* 20, 12 (2014), 2112–2121. <https://doi.org/10.1109/TVCG.2014.2346293>
 - [41] Zhenpeng Zhao, Sriram Karthik Badam, Senthil Chandrasegaran, Deok Gun Park, Niklas L.E. Elmqvist, Lorraine Kisselburgh, and Karthik Ramani. 2014. skWiki: A Multimedia Sketching System for Collaborative Creativity. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*. ACM, 1235–1244. <https://doi.org/10.1145/2556288.2557394>