

Textoshop: Interactions Inspired by Drawing Software to Facilitate Text Editing

Damien Masson
University of Toronto
Department of Computer Science
Toronto, Canada
damien.masson@umontreal.ca

Young-Ho Kim
NAVER AI Lab
Seongnam, Republic of Korea
yghokim@younghokim.net

Fanny Chevalier
University of Toronto
Department of Computer Science and
Statistical Sciences
Toronto, Canada
fanny@cs.toronto.edu

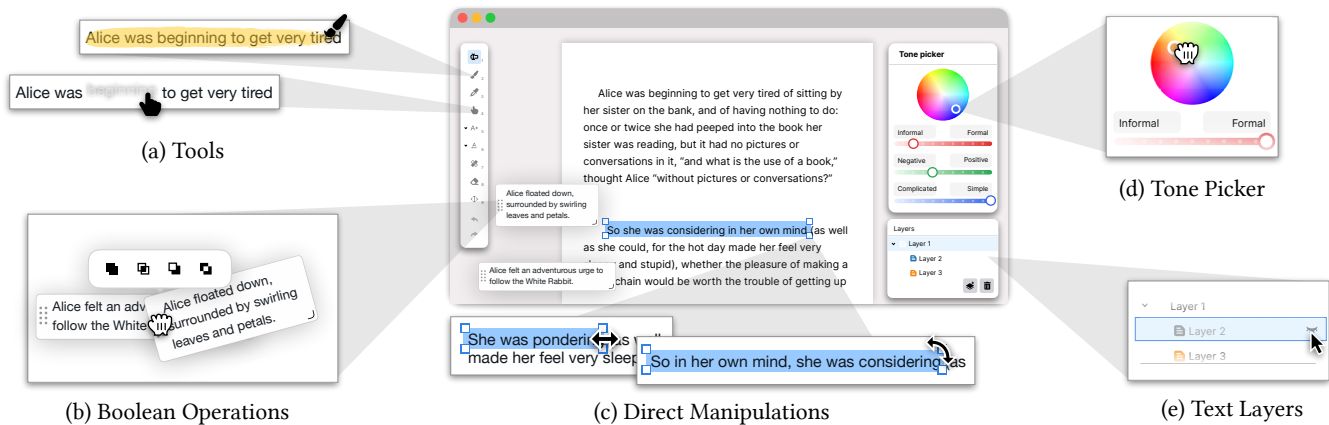


Figure 1: Textoshop uses interactions inspired by drawing software to help edit text: (a) text can be revised using tools such as the smudge and tone brush; (b) text passages can be merged using boolean operations; (c) text can be shortened and reordered through direct manipulations; (d) tones can be explored in a tone picker; and (e) text can be organized in layers

ABSTRACT

We explore how interactions inspired by drawing software can help edit text. Making an analogy between visual and text editing, we consider words as pixels, sentences as regions, and tones as colours. For instance, direct manipulations move, shorten, expand, and reorder text; tools change number, tense, and grammar; colours map to tones explored along three dimensions in a tone picker; and layers help organize and version text. This analogy also leads to new workflows, such as boolean operations on text fragments to construct more elaborated text. A study shows participants were more successful at editing text and preferred using the proposed interface over existing solutions. Broadly, our work highlights the potential of interaction analogies to rethink existing workflows, while capitalizing on familiar features.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI '25, April 26-May 1, 2025, Yokohama, Japan

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-1394-1/25/04...\$15.00

<https://doi.org/10.1145/3706598.3713862>

CCS CONCEPTS

• **Human-centered computing** → **Interactive systems and tools; Graphical user interfaces.**

KEYWORDS

writing, interface metaphors, drawing interaction, LLM, AI

ACM Reference Format:

Damien Masson, Young-Ho Kim, and Fanny Chevalier. 2025. Textoshop: Interactions Inspired by Drawing Software to Facilitate Text Editing. In *CHI Conference on Human Factors in Computing Systems (CHI '25)*, April 26-May 1, 2025, Yokohama, Japan. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3706598.3713862>

1 INTRODUCTION

The processes involved in writing share many similarities with those involved in image editing: words form sentences to tell a story, much like pixels come together to create regions that form a complete picture. Just as writers plan and refine a story by outlining the narrative and rearranging the story flow [26], visual artists plan and refine a graphical composition by laying out visual elements on a canvas, rearranging them, sometimes organizing them in layers [73]. Just as writers combine arguments into sentences and paragraphs [56], visual artists combine shapes into complex drawings [77]. And just as writers revise the language and iterate on the

wording and structure [24], visual artists revise the aesthetics and iterate on the colours and textures [73].

Despite similarities in *process*, the interaction paradigm to support digital writing differs substantially from the powerful ways to manipulate visual content. For example, in image editing and vector graphics software like Adobe Photoshop and Illustrator, artists can plan by directly manipulating elements of the image, moving them around and outside the canvas, and organizing them in independent layers. They can compose complex geometric structures by uniting, intersecting, or subtracting simple geometries. And they can rapidly explore nuances by navigating the space of colours in a colour picker. Even a task as simple as resizing an element, which drawing software support by dragging the selection box handles, will require careful rewriting with a text editor—think about the amount of effort involved in wordsmithing a paragraph to make it exactly fit within the space left on a page.

While large language models (LLMs) can help with writing tasks, how we interact with them rarely achieves the simplicity and expressiveness of drawing software. For example, LLMs must be prompted using relatively slow text input, preventing quick and precise iterations. These prompts are often verbose, and users may struggle to convey their intent using only words [18, 48]. In contrast, drawing software rely on direct manipulation of graphical elements. They also offer tools for smart edits. For example, the colour picker allows for exploring the space of solutions and expressing intents that are difficult to formulate in words. We see an opportunity to learn from this interaction paradigm to rethink the digital writing workflow.

To better support writers¹, we propose to adapt solutions used in drawing software to help with recurring writing tasks. Specifically, we aim to (1) capitalize on solutions that have been proven to work well and (2) offer solutions most users are already familiar with. Loosely, we view *words* as *pixels*, *sentences* as *regions*, and *tones* as *colours*. This analogy resulted in the design of Textoshop (Figure 1), a text editing software that supports meaning-preserving changes through direct manipulation to move, resize and restructure text selections; a layer management system to organize, display, and edit text in independent views; a sentence composition feature to construct sentences or paragraphs by combining, intersecting, or subtracting ideas or phrases; a tone picker to explore different tones when editing text; and tools including an “eraser” to intelligently remove part of the text, a “tone eyedropper” to copy the tone of a passage, and a “smudge” tool to paraphrase and smoothen transitions between sentences.

In a user study with 12 lay individuals who frequently engage in writing activities, Textoshop was compared to existing solutions when accomplishing text editing tasks such as organizing and integrating sentence fragments, changing the tone, using a template, and summarizing and expanding a passage. In a second part, participants used Textoshop in a free-form exploratory task. Results show that participants were more successful and preferred using Textoshop over alternatives. Participants appreciated the design of Textoshop and understood the analogy, allowing them to transfer some of their knowledge of drawing software to use the tool.

¹By “writers”, we do not necessarily mean professional writers but any individual who engages in writing activities, no matter their expertise

Overall, our work promotes leveraging interaction analogies to reimagine existing workflows.

2 BACKGROUND AND RELATED WORK

Traditional text editors such as Microsoft Word offer limited support for editing the text. We first review the issues this creates before discussing solutions proposed in the literature. We then detail the rationale for taking inspiration from drawing software².

2.1 Issues with Traditional Ways of Editing Text

First, *organization capabilities are limited when relying solely on the metaphor of sheets of paper*. Writers go through different stages in a non-linear order [25]. For example, they often alternate between generating ideas, constructing sentences, and revising these sentences [14]. In fact, it is often recommended for writers to organize ideas into spatial representations [26, 67]. Yet, text editors force organization in lines on pages, often leading to messy drafts mixing ideas, outlines, revisions, and discarded alternatives. This often results in a poor organization limiting writers’ exploration [62].

Second, *exploring the space of possible text variations is difficult*. A sentence can be rewritten in many ways by varying tone, word order, or length. In fact, rewriting is often considered “the essence of writing” [82]. Yet, most text editors offer little support to help revise language other than a thesaurus to find synonyms. Some experimental tools leverage LLMs to revise text, but it remains challenging for writers to express their intent with fine granularity [40].

Third, *reorganizing ideas already expressed is tedious* because ideas are difficult to isolate and text fragments are difficult to combine. Once ideas are expressed into sentences and paragraphs, moving these ideas requires rewriting most of the sentence or paragraph to guarantee consistency [41]. Besides making the exploration of alternatives difficult, it may also cause writer’s block as it reinforces this view of an immutable text structure [64].

Fourth, *managing and rendering different versions of a text is tedious*. Most text editors consider text as a linear sequence of characters with no structure other than chapters, sections, and paragraphs. Besides, edits are typically destructive, with only support for sequentially reversing the most recent changes incrementally with “undo”, or lookup in document histories that are tedious to browse [16]. This means that finding specific content and storing alternatives might be challenging. Writers may need to split their writing into different files or keep different copies of their text for each version, which is tedious to manage [31].

2.2 Novel Ways of Editing Text

Previous work has proposed solutions to some of the above issues, often using generative AI. Below, we describe these solutions and categorize them based on the problem they address.

2.2.1 Writing Using Non-linear Text Editors. As early as the development of the World Wide Web, Bolter and Joyce [8] proposed a writing tool supporting interactive fiction writing using a structural editor with a diagrammatic view. This spurred research into

²We use the term “drawing software” to refer to both raster (e.g., Adobe Photoshop) and vector (e.g., Inkscape and Adobe Illustrator) graphics editors.

the development of hypertext authoring tools [43] with which authors write part of the story in an order that might differ from how readers will read it. While these tools were primarily designed to create interactive stories, the idea of being able to write in any order inspired tools to author more traditional and static text. For example, the commercial software Scrivener [65] allows writers to compose text in any order by relying on different text areas that can be reorganized to rearrange the final story.

Textoshop also supports non-linear writing: text can be either positioned following lines on a page or freely arranged in 2D in and off the page. The margins can store mood boards, alternatives, and outlines akin to how digital artists leverage the space outside the canvas to preserve elements for reference [32]. Additionally, Textoshop supports layers of text can be edited independently and combined at will to create different composite stories.

2.2.2 Revising and Exploring Text Variations. While LLMs can generate text variations, steering the generation remains challenging. In its simplest form, previous work has proposed to select a passage and then prompt the model [33, 48] or click a button [62] to generate text alternatives. However, this interaction gives little freedom to control the result. Instead, other work proposed to generate not one but thousands of text variations in the hope that one of them might be satisfactory. For example, Gero et al. [27] explored different ways to present and query these different versions. Similarly, Luminare [72] arranges the different variations along specific dimensions, making it easier for users to find a variation of interest.

Our work differs in that the space of possible variations is explored through direct manipulation. If a sentence is close to satisfactory, users can explore variations of different lengths or word orders through the resizing and rotating features. They can also move a slider in a two-dimensional space to explore tone nuances, much like visual artists explore colour nuances in a colour picker.

2.2.3 Reordering and Combining Text Fragments. New writing tools increasingly “*shift the writer’s role towards editorial decisions*” and manipulating text fragments [11]. For example, Langsmith [36] can suggest full sentences from incomplete and rough sentence fragments. Similarly, Rambler [44] allows refining and combining dictated train-of-thoughts into full drafts. Other work proposed to manipulate summary representations such as paragraph summaries [22] and visual outlines [80] to revise the flow of ideas.

We go beyond simple merging of text passages by exploring novel ways to combine text inspired by boolean operations on shapes, such as subtracting to remove the information in one passage from another, intersecting to preserve only shared information between two passages, and excluding to remove redundant information redundant between two passages.

2.2.4 Text Versioning. Most text editors have some form of version control, albeit it is often limited. For example, Microsoft Word can track changes made to a document, and Google Docs and Scrivener [65] support storing, restoring, and comparing snapshots. More advanced versioning systems include ABScribe [62], which allows storing multiple text variations at specific text locations.

Our work explores the use of text layers that can also store modifications at specific text locations. Further, these layers can also store spatially disjoint but semantically related edits. This

allows supporting new workflows where writers store alternatives and hide or show layers to try out different versions. And this also supports new use cases such as templating.

2.3 Designing Interactions Based on Analogies

At its core, our approach repurposes interactions from one domain (in our case, image editing) to another domain (i.e. text editing). This capitalizes on features familiar to users and might inspire new workflows and new ways to think about the application domain.

Designing software interactions based on such analogies (also referred to as interface metaphors [51]) can lead to easily learned and expressive interactions [63, ch 2.4]. For example, the direct manipulation interaction paradigm is all about manipulating elements in software as if they were physical objects that can be moved, resized, or distorted [35, 68]. Instrumental interaction extends direct manipulation by noting that humans leverage tools when the finger or the hand is not enough [5, 6]. In the software domain, these tools are interaction instruments that operate on domain objects.

Drawing software combine a canvas analogy with ideas from direct and instrumental interaction. For example, Adobe Photoshop allows elements to be moved, removed, resized, and rotated through rapid and reversible physical actions that mimic how one interacts with physical objects in the real world, just like advocated by direct manipulation. And interaction instruments simulate paint brushes, pencils, or erasers. Other interaction instruments include the colour picker, which is analogous to a paint palette, and layers, which are analogous to transparent sheets [1, 7] and *reify* [6] the concept of independent manipulation.

The success of drawing software has been inspirational for other software. For example, it is common to find the concept of brushes, colour pickers, and layers adopted by software operating in different domains, such as presentation and video editing. And some of the digital drawing concepts have been adapted successfully, such as using a paintbrush to control the fortune of characters in a story [20], mixing prompts just like colours can be mixed to convey visual styles difficult to express with words [18], or repurposing raster image tools to manipulate colour histograms [17]. Beyond offering more expressive interactions, these analogies help people reason about unfamiliar software by leveraging their accumulated interaction knowledge about digital tools [61].

In contrast, text editors evolved from command-line word processors [30, 46] and use the analogy of the typewriter and sheet of paper [42]. In this work, we explore how interaction inspired by the graphical domain could open up new possibilities for writing.

3 TEXTOSHOP

We reimagine the design of text editors in Textoshop, a text editor that incorporates features and workflows borrowed from digital graphics editing. Textoshop employs a digital canvas instead of the sheet of paper metaphor. This means that text fragments are not seen as ink on a page but more like pieces of a collage that can be rearranged, resized, and modified at any time. Much like digital drawing software, Textoshop also offers smart tools to manipulate and modify these text fragments.

We designed Textoshop by reviewing features found in most raster or vector image editing software and that computer users

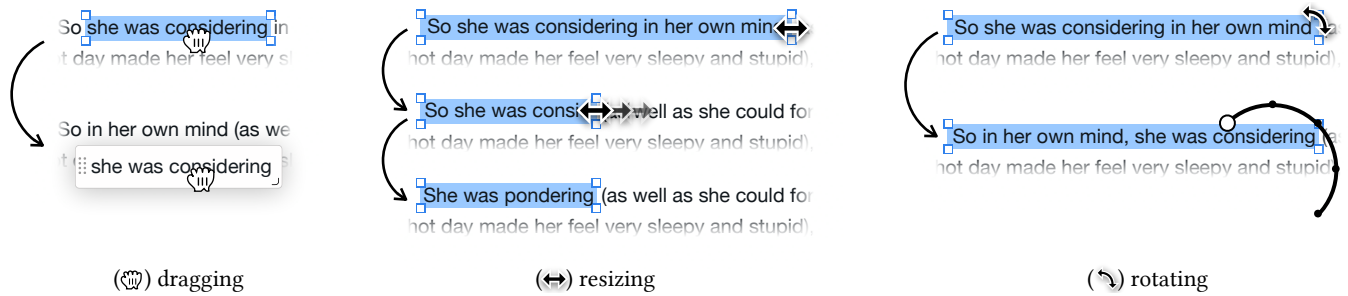


Figure 2: Direct manipulations on text selections edit the text: (☞) dragging moves the selection anywhere on the canvas; (↔) resizing shortens or expands the text; (↻) rotating changes the order of the words.

are likely familiar with, such as colour pickers, drawing tools, and layers. We then mapped these features to novel text editing capabilities, favouring fast interactions for frequent text editing tasks as identified in the taxonomy from Faigley and Witte [24]. Below, we categorize these features based on this taxonomy, which differentiates between formal changes (e.g., spelling, tense, number), meaning-preserving changes (e.g., substituting a word with a synonym, reordering words in a sentence), and micro/macro changes.

3.1 Dragging, Resizing, Rotating, and Splitting

Direct interactions are typically reserved for frequent operations. In Textoshop, we map these interactions to meaning-preserving changes whose impact is limited to the selection, such as additions, deletions, permutations, distributions, and consolidations [24].

☞ Dragging – Freely moves a text selection in and off the page through a drag-and-drop interaction (Figure 2.☞). Once dragged, the selected text is freed from the positioning constraints of the page and can be placed anywhere on or around the page on the canvas, supporting organizations similar to drawing software.

↔ Resizing – Lengthens or shortens the selected text to fit within the specified space. The interaction is triggered by starting a dragging motion from the right edge of the text (Figure 2.↔). The modification is meaning-preserving. As such, shortening often removes words, while lengthening adds words non-essential to the core meaning of the text. Large expansions of a text selection might result in added descriptions and new sentences.

↻ Rotating – Reorders the words in the selected text. The interaction is activated through a dragging motion starting on one of the four corner handles. The further the rotation thumb is moved from the center, the further the words are moved away from their original position (Figure 2.↻). For example, rotating “The ball was kicked by Lisa” reorders the words by using the active instead of the passive voice: “Lisa kicked the ball”.

⌘ Combining and ⇄ Splitting Sentences – Splits a sentence into multiple, or consolidates multiple sentences into one. The action is in a contextual menu opened by right-click a text selection.

3.2 Tools

Tools extend direct manipulation capabilities by activating modes that change the meaning of direct interactions [5, 6]. Textoshop

maps tools to modifications that are meaning-preserving but whose impact can go beyond the selection.

🖌 Tone Brush – Changes the tone of the selected text using the currently selected tone (in the tone picker, section 3.3). For example, brushing the text “Alice was beginning to get very tired” with an informal tone would result in “Alice was *totally wiped out*”.

👁 Tone Eyedropper – Copies the tone of the selected text and update the tone picker (section 3.3). For example, eyedropping the text “Alice was utterly exhausted” updates the tone picker to a tone mostly negative, slightly formal, and complicated. The copied tone can then be applied to another part of the text using the tone brush.

👉 Smudge – Regenerates the selected text. For example, smudging the word “beginning” modifies it to “starting”. Smudging can also help connect or change the transition between two sentences by smudging the end of a sentence and the start of the next one.

A- Singularizer and A+ Pluralizer – Singularizes or pluralizes the selection and updates the rest of the text to fix number agreements. For example, pluralizing “She” in the sentence “She was tired of sitting on her own” results in “*They were* tired of sitting on *their* own” where the verb and pronouns are also pluralized.

⏪ Past, ⏩ Present, and ⏭ Future Tense Changer – Changes the tense of the selection and updates the rest of the text to fix tense agreements. For example, applying the future tense to “was” in “Alice was beginning to tire” results in “Alice *will begin* to tire”.

🔧 Repair – Fixes the grammar, punctuation, spelling, and add necessary words to form valid sentences. For example, repairing “Alice vry tire” results in “Alice is very tired”.

🗑 Eraser – Removes the selected text and updates the rest of the text accordingly. Erasing makes sure the resulting text remains grammatically correct. Thus it might restructure the sentence and slightly modify other parts. For example, erasing “beginning” in the sentence “Alice was beginning to get very tired” results in “Alice was getting very tired”.

🗨 Prompt – Applies a prompt on the selected text. After a text selection, a text input box opens to allow the prompt to be entered. A large language model executes the prompt and the result is immediately integrated into the text.

3.3 Tone Picker

Textoshop adapts traditional colour pickers to allow navigating the space of tones and support meaning-preserving permutations. Essentially, tones are mapped to colours, allowing traditional colour space representations to visually represent the space of tones. Textoshop implements the RGB colour space. A tone in Textoshop is defined along three dimensions (Figure 3). We choose “formality”, “sentiment”, and “complexity” and map them to red, green, and blue based on word-colour associations [39]. Blending a mix of colours results in the corresponding tone, for instance, a “purple” tone means formal (red) and simple (blue), and mostly negative (no green), see Figure 3. These tones were selected as defaults because of their orthogonality and their use in previous work on manipulating tones [34, 38, 54]. These dimensions can be edited to explore other tone spaces, if needed. Further, we note that depending on user preferences and tasks, other colour spaces could be more suited.

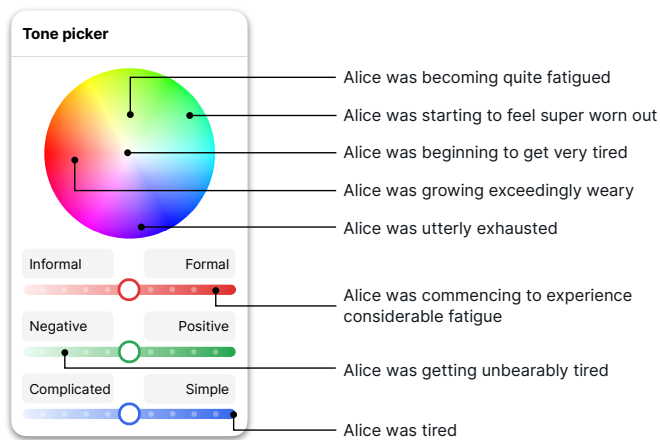


Figure 3: The tone picker allows exploring tone nuances of a sentence. Here, the selected sentence is “Alice was beginning to get very tired” and variations are explored by moving the thumb in the tone wheel and by changing the tone sliders.

Tone Sliders – Changes the tone along one specific tonality dimension. The position of the thumb in the tone wheel is updated accordingly. Each of the three sliders ranges from 0 to 10, allowing to represent $11^3 = 1331$ different tones.

Tone Wheel – Changes the tone along all three dimensions through one dragging motion in the colour wheel. Tone sliders are updated accordingly. To help users understand the tone wheel, arrows are overlaid, indicating the direction of the strongest change for each of the three tonality dimensions.

Exploring Other Tones – Tones in the colour picker are input text boxes that can be edited. Users can change the tones to explore new tone spaces, such as “humour” or “Shakespearean”, or extend a scale to have finer granularity, such as going from “formal” to “very formal” to offer 10 levels of formality.

Changing the Tone After or Before the Selection – The tone picker supports both noun-verb and verb-noun command constructions depending on the order of actions [59, ch 3-3]. In a noun-verb

construction, the text is selected first (noun), and following changes to the tone (verb) immediately update the passage, to rapidly explore different tone variations. In a verb-noun construction, the tone is selected beforehand (verb) and applied to future text selections (nouns) using the tone brush, to apply the same tone rapidly.

3.4 Boolean Operations on Sentences

Overlapping text can be merged in different ways to support micro and macro text changes. These merges are based on the ideas expressed in the text. Textoshop supports the following five methods to merge two pieces of text directly inspired by boolean operations on shapes (Figure 4).

■ **Unite** – Merges the two passages so that the resulting text is grammatically correct and includes the ideas from both passages. For example, merging “the dog is playing” with “the cat is playing” results in “the dog and the cat are playing”.

⊞ **Intersect** – Forms a grammatically correct text that expresses the ideas shared by both text passages. For example, intersecting “the cat and the dog are playing in the yard” and “the dog is chasing the cat in the yard” results in “the dog is in the yard”.

⊖ **Subtract** – Removes the ideas expressed in the first passage from the second passage. For example, subtracting “it’s raining today” from “she wears her coat because of the wet weather” results in “she wears her coat”.

⊘ **Exclude** – Removes the ideas that are shared in both passages to preserve only the non-redundant ideas. For example, excluding “Tom likes both chocolate and vanilla ice cream” from “Tom enjoys chocolate ice cream” results in “Tom likes vanilla ice cream”.

↓ **Insert** – Naively inserts the text passage at the exact caret position where the text was dropped, similar to pasting the text in a regular word processor. This allows converting a freely moving text selection back into text on the page.

3.5 Text Layers

Layers were originally introduced to drawing software to support non-destructive editing and allow picture parts to be edited independently [1]. These layers were inspired from the transparent plastic sheets used by the animation industry to compose scenes. But they rapidly became more powerful [37] and even essential to support the organization and versioning of drawings [73]. Writers share similar needs in that they want more organizational support [65] and non-destructive edits that can be retrieved and reversed [62].

In Textoshop, we propose to rely on the familiar layer metaphor. Text from documents can be organized in layers. These layers are similar to layers from drawing software in that they can be viewed as transparent sheets stacked on top of each other. However, unlike drawing layers, text layers store modifications that are anchored in the text from layers underneath it (Figure 5). The reason is that, unlike images, text is positioned relative to other text (e.g., a sentence is positioned after another sentence, but it makes little sense for it to be placed at specific coordinates). This means that the position of the text in a layer is relative to the text of the preceding layers. A modification to any of the preceding layers might reflow the text of the following layers. And layers can hide the text below them.

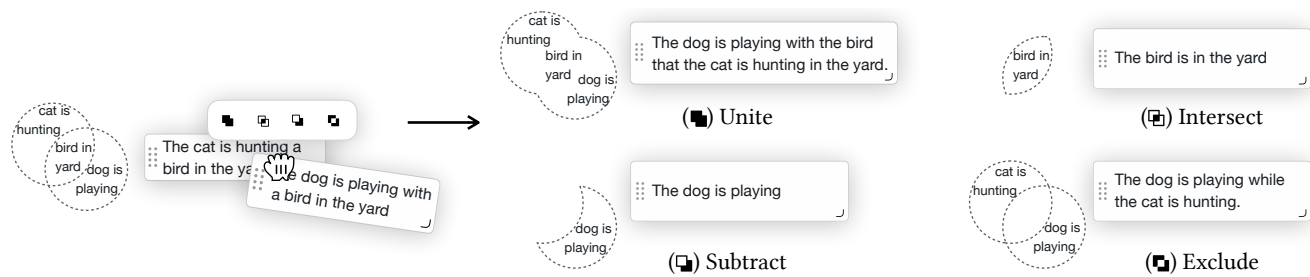


Figure 4: Dragging text over a passage allows merging the two passages in different ways: (U) “Unite” merges the two passages; (I) “Intersect” preserves only the information shared by both passages; (S) “Subtract” removes the information in one passage from the other; and (E) “Exclude” removes the information shared by both passages.

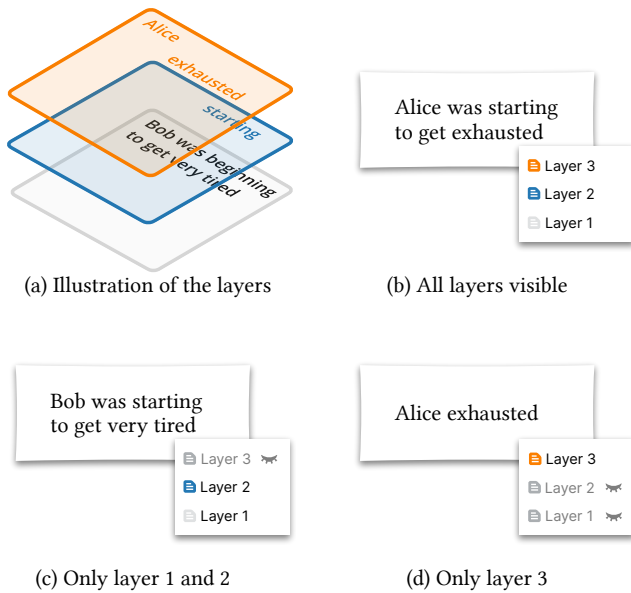


Figure 5: Text layers are stacked on top of each other: (a) each layer may store text anchored at different text locations; (b) layers at the top might hide the text underneath; (c) layers can be hidden; and (d) text is reflowed to avoid gaps

✚ **Adding a Layer** – Adds an empty layer at the top of the stack. Then, if this layer is selected, all new text modifications are stored in that layer. For example, text from the preceding layers can be modified or removed, and new text can be added, in which case it will be anchored within the text of the preceding layers.

👁 **Hiding and** 🗄 **Showing a Layer** – Hides/shows the modifications stored in a layer. For example, if a layer modified the text, hiding this layer restores the original text. And if a layer added some text, hiding the layer removes the added text.

🔄 **Reordering Layers** – Changes the order in which layers are shown. Reordering is done by a drag-and-drop interaction where the final position of the layer is decided by where it is dropped in the list. The order of layers decides what is visible. For example, the top

layer (i.e., the first in the list) is always visible because it appears above all other layers.

3.6 Implementation

The Textoshop interface emulates the look and feel of traditional drawing software, except the main canvas is replaced by a text input field (Figure 1). Under the hood, it uses a large language model as well as more traditional natural language processing techniques.

Textoshop is implemented in TypeScript using React [60] for the interface, Slate.js [71] for the text editor, NextUI [52] for the graphical components, and the OpenAI library [55] to execute prompts using OpenAI’s “gpt-4o” model. A live demo and the source code are available³.

3.6.1 Highlighting Text Changes. To help identify text changes, we added an animation to highlight character additions and deletions. The animation follows the guidelines from the transitioning technique between text revisions proposed by Chevalier et al. [16]: deleted characters appear first, in red, and progressively shrink and fade out. Added characters appear second, in green, and progressively grow and fade in until matching the size of the rest of the text. The whole animation lasts about one second.

3.6.2 Preserving Semantic Meaning. Certain modifications, such as resizing or relocating text, might result in text that is grammatically correct but with a different semantical meaning. We implemented several mechanisms to mitigate the loss of semantic meaning. For example, when resizing, the original sentence is saved and always used for subsequent prompts sent to the model. This means that even if the sentence is reduced to a couple of words, expanding it again will be based on the original longer sentence. Similarly, when relocating a passage, the model is prompted with the sentences surrounding the final location so that the sentence is adjusted to blend into its new location.

3.6.3 Prompting Strategies and Heuristics. Our intent with the implementation was to have robust, fast, and reliable tools to allow testing the concept behind Textoshop. While the implementation could perhaps be improved, we currently rely on engineered prompts and heuristics. For example, tools operating on a selection were optimized to only send the selection to the model and reintegrate the results by making it match the formatting of the original text.

³Live demo and code: <https://github.com/m-damien/Textoshop>

Another example is the resizing tool, which first splits the selection into sentences and prompts the model to shorten or expand each sentence in parallel. In our tests, GPT-4o was unreliable when prompted to resize a text by a certain amount of words. To boost the reliability of the resizing, the model was prompted to resize the same text eight times using a slightly different number of words. Then, the system calculated all possible combinations of the resized parts and selected the one closest to the target length.

4 USER STUDY

We ran a two-part user study to learn about the usability and usefulness of our approach. Specifically, we tried to answer the following questions: How do people use the features of Textoshop to accomplish typical writing tasks? How does their knowledge of drawing software help them understand Textoshop? Is Textoshop easier to use than existing solutions?

The first part of the study allowed us to isolate the usability of different features in a controlled setting. Each participant accomplished a series of tasks using either Textoshop or a baseline interface. For the baseline, given that no existing solutions are directly comparable to Textoshop in terms of features, we opted for an interface composed of a traditional text editor and an integrated replica of the ChatGPT interface, which participants could supplement with any external software as they saw fit.

The second part of the study was more exploratory to observe how participants use the tool as a whole in a more holistic and unguided way. Participants used the fully-featured version of Textoshop on their own text with no other instruction than to explore the tool and its features to accomplish their text editing goals.

4.1 Participants

We recruited 12 participants from our local community (age range: 23 to 35, Mdn=26; 7 self-identified as female and 5 as male). On a 5-point scale from 1-never to 5-very frequently, participants reported how often they engaged in writing activities (Mdn=5, SD=.7), how often they use generative AI to help them during writing activities (Mdn=2.5, SD=1.5), and how often they use drawing software (Mdn=3, SD=1). Most of their writing was academic (12), professional (11), and creative (2). On a 5-point scale from 1-not at all familiar to 5-extremely familiar, they also reported their familiarity with specific drawing software features such as colour pickers (Mdn=5, SD=.8), drawing tools (Mdn=4, SD=.7), layers (Mdn=4, SD=1.1), and boolean operations on shapes (Mdn=4, SD=.6).

4.2 Apparatus

We ran the study remotely using the Zoom conference software. Participants used their personal computers and used a modern web browser to open the experiment website. The whole session was audio and screen recorded, and the study software logged all interactions with the system (e.g., clicks, features used, answers to questions). Sessions were about an hour long. Participants received a CAD\$20 gift card in appreciation for their time. Our institutional ethics board approved the study.

4.3 Procedure

Each participant used two INTERFACES (Textoshop and baseline). For each INTERFACE, the participant completed 4 TASKS (shortening and expanding a story, organizing and integrating text fragments, changing the tone of dialogue, and using a template to write emails).

Introduction (~2min) – The participant completed the consent form and demographics questionnaire and was told they would use two different interfaces to accomplish a series of text editing tasks.

Video Tutorial (2 INTERFACES × 4 TASKS × 20sec) – Before each task, the participant watched a 20-second video demonstrating the feature they are about to use. The videos demonstrate features on abstract tasks, akin to the typical tool previews found in modern software [28]. The video for the baseline condition demonstrates the use of our replica of ChatGPT, and the participant was allowed to skip it after having watched it once. No video was provided for the external software the participant might have chosen to use.

Text Editing Tasks (2 INTERFACES × 4 TASKS × 4min) – After watching the video mini-tutorial, the participant was asked to accomplish a text editing task using a given INTERFACE. In the baseline condition, besides the provided integrated version of ChatGPT, participants were encouraged to use external software of their choosing to accomplish the task if they felt it would be easier. Instructions for the task were shown in the bottom left corner and consisted of a short textual description. The participant had to manually indicate when done with the task by pressing a button. Otherwise, after 4 minutes of working on the task, the participant was asked to stop and move to the next task. The participant then had to indicate how successful they felt in accomplishing the task on a 5-point semantic differential scale from 1-unsuccessful to 5-successful.

Questionnaire About Interface Used (2 INTERFACES × 2min) – After completing the 4 tasks with an INTERFACE, the participant completed a system usability scale (SUS) questionnaire [10] and answered 5-point scale questions related to the difficulty of the tasks they just completed. The second and final questionnaire also included comparative scale questions asking the participant to report the interface they preferred to accomplish each task.

Free-form Exploration with Textoshop (~10min) – In a second part, the participant was invited to use the fully-featured version of Textoshop on their own text. Before the study, the participant had been instructed to bring a text they “would be interested in editing (about 4-5 paragraphs long)”. In case they did not bring text, they were also allowed to choose amongst a collection including Wikipedia articles and passages from novels. Because we meant for this part to be a free-form exploration, the participant received no instruction other than “to explore the tool and its features to accomplish their editing goals”. Given that some of the features of Textoshop had not been used nor demonstrated in the previous part, the participant could also ask for an explanation of these features.

Semi-structured Interview (~10min) – After finishing the free-form exploration, the participant engaged in a semi-structured interview with the experimenter. Questions were open-ended and related to their experience with the system, their preferences, and their impressions on the design of the tool and its usability.

4.4 Design and Tasks

The study followed a within-subject design. Interface order was counterbalanced. To avoid participants doing the same tasks with both interfaces, we designed two variations of each task using different but comparable texts. The order of the tasks followed a balanced Latin square, and the variation used was picked at random.

We designed four tasks such that they cover the issues with traditional ways of editing text (section 2.1) and involve as many of the features of Textoshop as possible. The tasks were designed to be realistic, require multiple steps, and ask of participants to express higher-level intents rather than simple editing operations. Rather than evaluating all tools, our main question was whether participants understand and successfully use features building on the drawing software analogy. The four tasks were as follows.

Shortening/Expanding a Summary – Given a summary of a well-known novel, the participant has to shorten and expand specific sentences and paragraphs. For example, they need to compress a paragraph to avoid the short last line (“orphan”) while preserving the meaning. Or they need to expand upon a specific sentence in another paragraph so that it covers two lines and fills the whole width of the page. With Textoshop, the participant can use direct manipulation to resize text selections. This task emulates the frequent need for writers to change the length of their text so that it fits within a spatial layout (e.g., posters and presentations) or remains below a character or page limit.

Changing the Tone of Dialogues – Given a dialogue between two fictional characters, the participant has to change the tone of the dialogue of one of these characters (e.g., make their dialogues more informal and positive) and then explore at least four variations of the other character’s last dialogue. With Textoshop, the participant can use the tone picker and tone-brush. This task was to test the tone picker metaphor. It emulates copyediting situations when the tone of a document has to be uniformized (e.g., integrating changes from different writers) and when the tone has to be explored precisely to match a certain intent (e.g., writing an email to different audiences).

Organizing and Integrating Text Fragments – Given a simple story and a bullet list with 6 text fragments, the participant has to sort the text fragments depending on the character they are related to and then integrate these fragments into the story. Some text fragments have to be merged into specific parts of the text, while others are concepts to remove from the story. In the Textoshop condition, participants can arrange the text fragments in and off the page and can also use boolean operations to unite, intersect, subtract, or exclude text fragments. This task was to test the metaphor of the canvas and floating elements, as well as boolean operations on text fragments. It emulates editing the flow of ideas and rapidly iterating on different ways to structure these ideas.

Using an Email Template to Answer Customers – Given the profile of two customers as well as the email template used by a fictional company to answer their customers, the participant has to generate an email in response to each customer. The template includes pre-written paragraphs, which should be selected if they apply to that specific customer, and has to be edited further to include the name of the customer and some specific information, such as amounts and durations. With Textoshop, layers are already created with the

text of the email template; to form the final emails, the participant has to show/hide different layers and apply minor edits to the layers. With the baseline interface, the email template lists all the possible phrasings to answer a customer organized in sections; to form the final emails, the participant has to copy and paste them in order. This task was to test the layer metaphor and participants’ mental model of layers when editing. It simulates having to produce different text versions simultaneously based on different audiences, such as email communication or educational resources.

4.5 Data Analysis

Considering ongoing debates about statistics in HCI [23, 29], we report both confidence intervals (CI) and p-values. However, we base our conclusions on the 95% confidence intervals (CI) on mean differences, which provide information about both the presence and size of effects. All confidence intervals are calculated using the studentized bootstrapping method because it is the most robust for within-subject study designs of about 12 participants [47, 81]. All p-values are calculated using a Wilcoxon signed-rank test because it is less likely to yield false positives compared to parametric alternatives when dealing with data whose distribution is unknown and possibly heavily skewed [9]. The bootstrapped CIs are calculated in Python using arch 5.5.0 [66] and using 10,000 replications. The Wilcoxon signed-rank tests are calculated using scipy 1.10.1 [74] with ties being discarded.

4.6 Results from the Text Editing Tasks

Participants completed the tasks without major difficulties. During the baseline condition, and despite being encouraged to use external tools, only P11 used Grammarly. Other participants used a mix of the provided replica of ChatGPT and manual editing.

4.6.1 Success rate. On a 5-point semantic differential scale from 1-unsuccessful to 5-successful, participants rated being more successful at accomplishing the tasks using Textoshop than baseline by 0.71 (95% CI [0.27, 1.17], $M=4.6$ vs $M=3.9$, $p=.015$).

4.6.2 Usability. On the system usability scale and for the given tasks, Textoshop received a 90 (typically considered “excellent” [3]) while the baseline received a 69 (typically considered “ok”). This corresponds to a difference of 21.25 (95% CI [12.5, 32.91], $p<0.001$).

4.6.3 Time and Prompts Used. Although productivity was not our focus, we provide indications of completion time and prompts used as a loose indicator of the tedium involved in completing a task. Participants completed tasks faster using Textoshop than baseline by 55s (95% CI [26s, 1min 26s], $M=2\text{min } 23\text{s}$ vs $M=3\text{min } 19\text{s}$, $p=0.003$). Ranked from largest to smallest effects: “Shortening/expanding a summary” by 1min 29s, “Using an email template” by 1min 19s, “Changing the tone of dialogues” by 41s, and “Organizing and integrating text fragments” by 13s. Additionally, in the baseline condition, participants used 2.7 prompts per task on average, and these prompts were about 269 characters long. Considering participants did not use prompts in the Textoshop condition, we discuss participants’ behaviours in section 5.1.

4.6.4 Questionnaires. Participants responded to all eight 5-point statements in favour of Textoshop on average. Figure 6 shows the

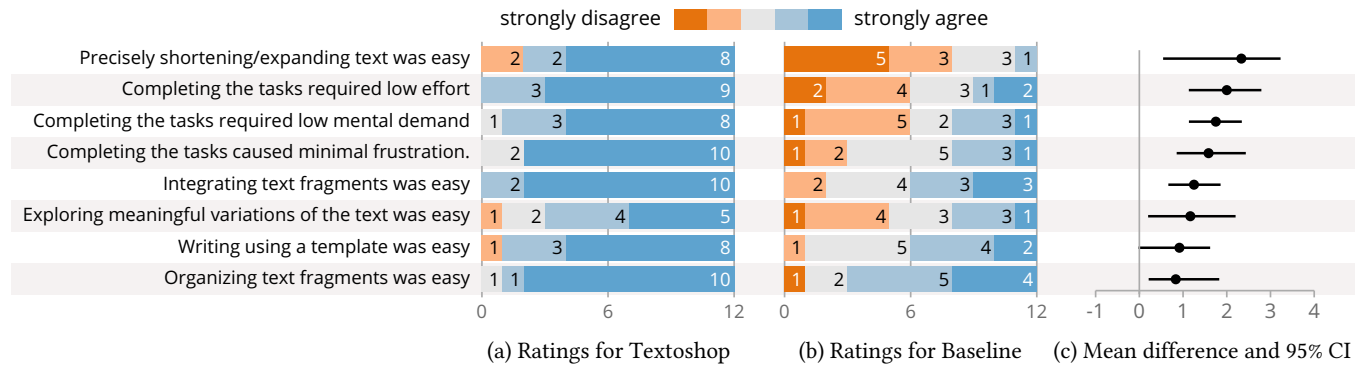


Figure 6: Participants’ response when rating the 5-point statements for (a) Textoshop and (b) baseline. (c) Dots are the mean differences of Textoshop compared to baseline. Bars are the 95% CIs calculated with the studentized bootstrap method.

breakdown of participants’ answers. Ranked from largest to smallest effects: “Precisely shortening/expanding text was easy” by 2.33 (95% CI [0.54, 3.23], $p=.007$), “Completing the tasks required low effort” by 2.0 (95% CI [1.13, 2.79], $p=.004$), “Completing the tasks required low mental demand” by 1.75 (95% CI [0.85, 2.44], $p=.003$), “Completing the tasks caused minimal frustration.” by 1.58 (95% CI [0.66, 2.44], $p=.007$), “Integrating text fragments was easy” by 1.25 (95% CI [0.66, 1.86], $p=.006$), “Exploring meaningful variations of the text was easy” by 1.17 (95% CI [0.2, 2.2], $p=.026$), “Writing using a template was easy” by 0.92 (95% CI [-0.02, 1.62], $p=.053$), “Organizing text fragments was easy” by .83 (95% CI [0.21, 1.83], $p=.033$).

4.6.5 Preferences. When asked about their preferred INTERFACE for each task, participants unanimously rated Textoshop as “preferred” or “strongly preferred” when “integrating text fragments”. All participants also preferred Textoshop when “precisely shortening/expanding text” except P10, who did not have a preference for either interface. For other tasks, while the majority preferred Textoshop, some participants expressed other preferences: when “writing using a template” P8 preferred the baseline interface; when “organizing text fragments” P3 preferred and P8 strongly preferred the baseline; and when “exploring meaningful variations of text” P7, P8, and P12 preferred and P3 strongly preferred the baseline. We discuss the participants’ reasons for these preferences in Section 5.

4.7 Results from the Free-form Exploration

In the free-form exploration, participants chose texts related to creative story writing (P2, P4, P5, P9), informative writing (P3, P10, P11, P12), academic writing (P1, P6, P7), and personal writing (P8). Figure 7 shows the sequence of features used by each participant. Below, we report some of the strategies and behaviours observed.

4.7.1 Iterating on a sentence by alternating tools. Participants used the tools to iterate on the wording of some specific sentences, much like they would iterate on the colours, texture, and organization of elements in drawing software. However, unlike visual artists who often apply colours and texture only after sketching [73], with Textoshop, tone exploration happened right from the beginning. Often, participants’ workflow involved alternating between trying different tones, shortening the sentence, and smudging some parts to be paraphrased. The order of these steps was not always the

same, but the tone picker was often used first, and then the other tools would allow editing with more precision. With the exception of P2, who frequently rotated and split the sentences, other tools were seldom used to rewrite a sentence.

4.7.2 Experimenting with alternative organizations of ideas. P5, P10, and P12 reorganized the ideas multiple times by taking a passage from one paragraph and integrating it into a different paragraph or sentence using the boolean operations. P6 did the same but created their own text fragments to add an idea to an existing paragraph. Finally, P1 went one step further and deconstructed a whole paragraph into passages. Then, they organized these passages in the margin to try a different organization of the arguments and then reintegrated the passages into the main text using the “insert” and “unite” features. This suggests Textoshop encouraged writers to work on the flow of ideas even after the writing is fully-fledged into sentences and paragraphs. This is in contrast with traditional writing, where restructuring mostly happens in the planning phase, when the draft is only an outline [26].

4.7.3 Using layers to identify changes, store AI modifications, or add structure. Layers were used in multiple ways, most differing from how they are used in drawing software. For example, P1, P8, and P9 used one layer to store all the modifications made using the tools, allowing them to preserve the integrity of their original text and easily identify AI changes by switching the visibility of the layer. P3 used layers to store alternative versions of paragraphs (in this case, storing a more informal version). Perhaps the closest use of text layers to drawing layers that we have observed is to use them to structure the text. For example, P5, P11, and P12 created layers to store the different paragraphs, passages, or main ideas. P7 used layers to store the examples mentioned in the text. And P6 used a layer to store the statistics reported in the text so that they could read and edit their results section while hiding the means or p-values reported in the text (by hiding the statistics layer). This differs from traditional writing where text cannot usually be teased apart in semantically meaningful ways. It is akin to the organization proposed by systems like Scrivener [65] except with a much finer granularity than just chapters and sections.

software.. so... it's easier for me". In fact, P1 even commented on how they could guess how a tool would work before using it: "I didn't really use the eyedropper, but I saw it and I could figure out what it'd do: pick up the tone of whatever sentence I placed it on".

5.1 Observations and Participants' Comments

5.1.1 The drawing analogy was mostly successful. Good interface analogies "help establish user expectations and encourage predictions" [51]. We observed such expectations and predictions from our participants, as they mentioned how Textoshop allowed a trial-and-error exploration [49] similar to the exploration they were used to in drawing software. P6 mentioned "when trying to pick colours for my own drawing, I usually randomly just adjust the point on the palette and just kind of see what I can get. So I adopted the same strategy when I was trying tones" suggesting part of the success of the colour picker analogy is that it matched how one would use it in drawing software. However, one danger with interface analogies is the "overly literal translation of existing bad designs" [58]. We observed one instance of such phenomenon regarding the boolean operations, P1 explained "Even now with photo and image editing [software], 'unite' is pretty straightforward, but all the other ones [...] I often can't remember which one means what". Finally, many participants described the system as "fun" and P3 commented "I thought it was fun just to play around with it [...] it kind of sparked that same joy that I get when I mess around with image editing software". This relates to the observation from Carroll and Thomas [13] that an interface analogy that "is conducive to the desired emotional attitude of the user" and "provide the user with exciting metaphors for routine work" may enhance the intrinsic motivation of users in accomplishing their tasks.

5.1.2 Participants quickly learned how to use Textoshop, but it is unclear if they know when and how to integrate it into their workflow. The user study tested the first part of the learning curve of Textoshop, which corresponds to the time it takes for a user to reach a reasonable level [53]. Our results and observations suggest participants started high on the learning curve, possibly thanks to transferring their knowledge from drawing software [51, 57]. However, we only have anecdotal evidence about the shape of the second part of the learning curve, which corresponds to the transition from novice to expert. Specifically, we observed participants had no issues trying features and workflows and undoing them if necessary, which might help discover new features and expert workflows [49]. However, it remains unclear if people would actually integrate features of Textoshop into their workflow. For example, it is known that, despite knowing more efficient ways of accomplishing tasks, such as keyboard shortcuts, users might still prefer their inferior approach [4, 21]. Similarly, users and Textoshop might co-adapt such that features of Textoshop are appropriated and used in unexpected ways [45].

5.1.3 Identifying changes was sometimes challenging with Textoshop. Some participants noted how a main difference with drawing software was that it is not as immediate to perceive textual changes. P8 commented "in like Figma [...] it's immediately, visually obvious what changed. [...] It didn't feel like I had the same mobility here where I could like immediately understand what the changes and

the consequences were". In fact, some participants mentioned how they would prefer to have some time to review the changes before integrating them. P8 added "I really want to see things side-by-side". P6 explained that they would like a system to track the changes because the current animation of changes "disappeared so quickly before I can notice [the changes]". In Textoshop, changes can be tracked either by working on copies or using layers, but these comments seem to indicate that, unlike drawing software, this is not enough to help users identify textual changes.

5.1.4 Textoshop forced pointing interactions. Unlike traditional text editors that can be used using only the keyboard, Textoshop requires some pointing interaction to arrange text fragments, resize, rotate and manipulate tools. This was noted by both P7 and P8. P7 said "one difficulty is that [drawing software] is more heavily dependent on mouse or pen, so it's not about using the keyboard [...] one issue is switching between the two". P8 explained "I feel like text editing is inherently like a keyboard-driven task, and I felt like it was difficult to do that with [Textoshop]. [...] I think I just really know how to use the keyboard really quickly and fast".

5.1.5 Textoshop gave participants more control. Even though both Textoshop and the baseline used the same underlying model, participants frequently mentioned control as an advantage of Textoshop. For example, when referring to the tone picker, P4 explained "when I'm writing and I want a certain tone, it's nice to have a slider for that, because a lot of the times, text that ChatGPT generates is like weirdly wordy and formal. But then, when I ask make it casual, it makes it way too casual". When referring to the resizing feature, P11 explained "I don't want too much. The fact that I can just control like how much I want to elaborate on it, I think that was really cool and like, really useful".

5.2 Limitations

5.2.1 Textoshop might not be easy to learn for people with no experience with drawing software. While some participants reported not frequently using drawing software, all of them reported being familiar with drawing tools, colour pickers, layers, and boolean operations on shapes. We expect Textoshop to lose its advantage in terms of learnability when used by people with little to no knowledge of drawing software.

5.2.2 Due to the implementation of Textoshop, it inherits all the issues of using LLMs to edit text. While the implementation of Textoshop is not core to our contribution, it is important to acknowledge that Textoshop currently heavily relies on OpenAI's GPT-4o. It is known that LLMs have issues in terms of biases, privacy, and plagiarism [50, 76]. We also found the model to perform poorly on tasks such as resizing to a specific number of words. In our implementation, we workaroud these issues by regenerating multiple versions until one that matches the length is found. But it is often imperfect, which might frustrate users. Finally, some text transformations might result in imperfect, incongruent, or hallucinated text. These issues might be mitigated by using fine-tuned LLMs that take greater care in preserving semantic meaning and matching the style of the original text.

5.2.3 The drawing software analogy might confuse users. Using interface analogies might lead to mismatches between the source and the target domain. While some mismatches might be beneficial and lead to curiosity, others might cause problems and create invalid assumptions [51]. With Textoshop, this difficulty is coupled with the use of an LLM that might produce unexpected text, making the system difficult to model for the user [19]. While this did not seem to be an issue during the study, some participants did mention a mismatch between their mental model and the behaviour of the system regarding the layer system. For example, P3 mentioned not seeing text layers like stacked sheets but more like “*variables or conditions that you’re kind of swapping in and out*”. One possible explanation is given by Neale and Carroll [51] in that “mismatches occur with aspects of computer functionality when the target domain is unable to provide a direct real-world correlate”. Text layers have no real-world equivalent. Unlike physical transparent sheets, text layers essentially replace text and are positioned dynamically based on their anchor points in the text. Finally, text layers also led users to make invalid assumptions. For example, P8 expected features and were surprised not to find them “*In Figma [...] when I click again it goes down a layer rather than staying at the same level*”

5.3 Future Work

5.3.1 Repurposing other digital drawing features. We only scratched the surface of drawing features that could be repurposed for text. For example, layers could support different blending modes to either overwrite or track the changes and changing the opacity of the layer could control how the text is integrated (e.g., 100% overwrites, while 50% generates text with a meaning midway between the layer and what it overlaps). Other layer types could include “Adjustment layers” [1] that apply a transformation to the text (e.g., making the text more positive) while retaining the integrity of the text underneath. The rotation feature could allow specifying the centre of rotation (the word that should remain at the same location). Finally, text could be transformed via histograms showing text properties such as the number of adverbs, the sentiment of the text, or the length of the sentences. Then, writers could manipulate the histogram to transform the text.

5.3.2 Leveraging recent research in graphical editing to also benefit Textoshop. The design of drawing software is a well-researched topic. Much of this research could be adapted and implemented in Textoshop. For example, work on reifying [6] graphical attributes into manipulable objects [78] could be adapted so that properties such as tone, sentence length, and sentence rotation become first-class objects that can be easily manipulated, copied, reversed, and removed. Work on better colour pickers and colour theme authoring could similarly power better tone pickers where tones become objects that can be rearranged and interpolated [70]. Similarly, work on using histograms as another editing space could inform the design of text histograms that can query and edit the text [17].

5.3.3 Testing the impact of Textoshop over a longer period of time. While challenging due to the nondeterministic nature of LLMs that might confound the results, a longitudinal or diary study of writers using Textoshop over several months would tell us how people integrate the features in their workflows, how they co-adapt with

the system [45], and how it impacts the ideation, the quality of the result, and the productivity [19]. For example, the consequences of constantly switching between keyboard and pointing devices are unclear. Moving from the keyboard to the mouse takes about 0.4s [12] and might have a cognitive cost which might add up in the long run. However, we note that these switches may not be so frequent considering writing is not only about typing characters on a keyboard but also involves planning and reviewing, which do not leverage the keyboard [25]. Additionally, Buschek [11] observed a shift in how writing is done to look more like an assemblage of text snippets rather than manual writing. Such workflow is less reliant on the keyboard and might be better supported by Textoshop. Beyond productivity, a longitudinal study would also help study how Textoshop changes the writing workflow and the text that is produced. For example, the diverse ways participants used layers to structure their text is reminiscent of fluid documents [15] where the added structure and ways to reveal in-context information changed people’s reading behaviours [79]. While some features of Textoshop may not make writing easier nor faster, they might contribute to other aspects such as idea generation or quality of the craft [19].

5.3.4 Applying the drawing software analogy to other domains. Analogies used to design software often stay within the same domain: the digital canvas for drawing software and the blank sheet for text editors. Through this work, we show the potential of using analogies from different domains. We believe this strategy could inform the design of software from all domains. For example, code editors could also benefit from layers to structure the code. This would allow an organization with a finer granularity than files, and layers could be hidden to remove the code responsible for a specific feature. Algorithms could be merged, intersected or excluded to change their behaviours. And tools such as brushes could perform smart editing of code to refactor, add comments, or change the coding style. In fact, this idea was recently experimented with “Code Brushes” that can enhance the readability or fix bugs in code [75].

6 CONCLUSION

We presented Textoshop, a text editor that repurposes features from drawing software to offer new workflows and features that facilitate text editing. Results of a user study showed that participants were more successful and preferred Textoshop to accomplish text editing tasks. The study also highlighted how participants transferred their knowledge of drawing software to learn how to use Textoshop. Broadly, our work highlights the potential of software analogies to inform the design of software from any domain.

ACKNOWLEDGMENTS

This work was supported by NAVER corporation as a part of NAVER-Wattpad-University of Toronto research center and NSERC Discovery Grant RGPIN-2018-05072.

REFERENCES

- [1] Philip Andrews. 2006. *Adobe Photoshop Elements 5.0: A Visual Introduction to Digital Photography* (1st edition ed.). Routledge, Oxford ; Burlington, MA.
- [2] Kenneth C Arnold, April M Volzer, and Noah G Madrid. 2021. Generative Models Can Help Writers without Writing for Them. In *IUI Workshops*. ACM, Virtual Event.

- [3] Aaron Bangor, Philip T. Kortum, and James T. Miller. 2008. An Empirical Evaluation of the System Usability Scale. *International Journal of Human-Computer Interaction* 24, 6 (July 2008), 574–594. <https://doi.org/10.1080/10447310802205776>
- [4] Nikola Banovic, Fanny Chevalier, Tovi Grossman, and George Fitzmaurice. 2012. Triggering Triggers and Burying Barriers to Customizing Software. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. Association for Computing Machinery, New York, NY, USA, 2717–2726. <https://doi.org/10.1145/2207676.2208666>
- [5] Michel Beaudouin-Lafon. 2000. Instrumental Interaction: An Interaction Model for Designing Post-WIMP User Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '00)*. Association for Computing Machinery, New York, NY, USA, 446–453. <https://doi.org/10.1145/332040.332473>
- [6] Michel Beaudouin-Lafon and Wendy E. Mackay. 2000. Reification, Polymorphism and Reuse: Three Principles for Designing Visual Interfaces. In *Proceedings of the Working Conference on Advanced Visual Interfaces (AVI '00)*. Association for Computing Machinery, New York, NY, USA, 102–109. <https://doi.org/10.1145/345513.345267>
- [7] Matt Belge, Ishantha Lokuge, and David Rivers. 1993. Back to the Future: A Graphical Layering System Inspired by Transparent Paper. In *INTERACT '93 and CHI '93 Conference Companion on Human Factors in Computing Systems - CHI '93*. ACM Press, Amsterdam, The Netherlands, 129–130. <https://doi.org/10.1145/259964.260145>
- [8] Jay David Bolter and Michael Joyce. 1987. Hypertext and Creative Writing. In *Proceedings of the ACM Conference on Hypertext (HYPERTEXT '87)*. Association for Computing Machinery, New York, NY, USA, 41–50. <https://doi.org/10.1145/317426.317431>
- [9] Patrick D Bridge and Shlomo S Sawilowsky. 1999. Increasing Physicians' Awareness of the Impact of Statistics on Research Outcomes: Comparative Power of the t-Test and Wilcoxon Rank-Sum Test in Small Samples Applied Research. *Journal of Clinical Epidemiology* 52, 3 (March 1999), 229–235. [https://doi.org/10.1016/S0895-4356\(98\)00168-1](https://doi.org/10.1016/S0895-4356(98)00168-1)
- [10] John Brooke. 1995. SUS: A Quick and Dirty Usability Scale. *Usability Eval. Ind.* 189 (Nov. 1995), 7.
- [11] Daniel Buschek. 2024. Collage Is the New Writing: Exploring the Fragmentation of Text and User Interfaces in AI Tools. <https://doi.org/10.1145/3643834.3660681> arXiv:2405.17217 [cs]
- [12] Stuart K. Card, William K. English, and Betty J. Burr. 1978. Evaluation of Mouse, Rate-Controlled Isometric Joystick, Step Keys, and Text Keys for Text Selection on a CRT. *Ergonomics* 21, 8 (Aug. 1978), 601–613. <https://doi.org/10.1080/00140137808931762>
- [13] John M. Carroll and John C. Thomas. 1982. Metaphor and the Cognitive Representation of Computing Systems. *IEEE Transactions on Systems, Man, and Cybernetics* 12, 2 (1982), 107–116. <https://doi.org/10.1109/TSMC.1982.4308795>
- [14] Tuhin Chakrabarty, Vishakh Padmakumar, Faeze Brahmaan, and Smaranda Muresan. 2024. Creativity Support in the Age of Large Language Models: An Empirical Study Involving Emerging Writers. arXiv:2309.12570 [cs]
- [15] Bay-Wei Chang, John D Mackinlay, and Polle T Zellweger. 2000. Fluidly Revealing Information in Fluid Documents. In *Smart Graphics 2000 AAAI Spring Symposium*. AAAI, Palo Alto, California.
- [16] Fanny Chevalier, Pierre Dragicevic, Anastasia Bezerianos, and Jean-Daniel Fekete. 2010. Using Text Animated Transitions to Support Navigation in Document Histories. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, Atlanta Georgia USA, 683–692. <https://doi.org/10.1145/1753326.1753427>
- [17] Fanny Chevalier, Pierre Dragicevic, and Christophe Hurter. 2012. Histomages: Fully Synchronized Views for Image Editing. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology (UIST '12)*. Association for Computing Machinery, New York, NY, USA, 281–286. <https://doi.org/10.1145/2380116.2380152>
- [18] John Joon Young Chung and Eytan Adar. 2023. PromptPaint: Steering Text-to-Image Generation Through Paint Medium-like Interactions. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. ACM, San Francisco CA USA, 1–17. <https://doi.org/10.1145/3586183.3606777>
- [19] John Joon Young Chung, Shiqing He, and Eytan Adar. 2021. The Intersection of Users, Roles, Interactions, and Technologies in Creativity Support Tools. In *Designing Interactive Systems Conference 2021*. ACM, Virtual Event USA, 1817–1833. <https://doi.org/10.1145/3461778.3462050>
- [20] John Joon Young Chung, Wooseok Kim, Kang Min Yoo, Hwaran Lee, Eytan Adar, and Minsuk Chang. 2022. TaleBrush: Sketching Stories with Generative Pretrained Language Models. In *CHI Conference on Human Factors in Computing Systems*. ACM, New Orleans LA USA, 1–19. <https://doi.org/10.1145/3491102.3501819>
- [21] Andy Cockburn, Carl Gutwin, Joey Scarr, and Sylvain Malacria. 2015. Supporting Novice to Expert Transitions in User Interfaces. *Comput. Surveys* 47, 2 (Jan. 2015), 1–36. <https://doi.org/10.1145/2659796>
- [22] Hai Dang, Karim Benharrak, Florian Lehmann, and Daniel Buschek. 2022. Beyond Text Generation: Supporting Writers with Continuous Automatic Text Summaries. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. ACM, Bend OR USA, 1–13. <https://doi.org/10.1145/3526113.3545672>
- [23] Pierre Dragicevic. 2016. Fair Statistical Communication in HCI. In *Modern Statistical Methods for HCI*, Judy Robertson and Maurits Kaptein (Eds.). Springer International Publishing, Cham, 291–330. https://doi.org/10.1007/978-3-319-26633-6_13
- [24] Lester Faigley and Stephen Witte. 1981. Analyzing Revision. *College Composition and Communication* 32, 4 (1981), 400–414. <https://doi.org/10.2307/356602> jstor:356602
- [25] Linda Flower and John R. Hayes. 1981. A Cognitive Process Theory of Writing. *College Composition and Communication* 32, 4 (1981), 365–387. <https://doi.org/10.2307/356600> jstor:356600
- [26] Linda S. Flower and John R. Hayes. 1977. Problem-Solving Strategies and the Writing Process. *College English* 39, 4 (1977), 449–461. <https://doi.org/10.2307/375768> jstor:375768
- [27] Katy Ilonka Gero, Chelse Swoopes, Ziwei Gu, Jonathan K. Kummerfeld, and Elena L. Glassman. 2024. Supporting Sensemaking of Large Language Model Outputs at Scale. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–21. <https://doi.org/10.1145/3613904.3642139>
- [28] Tovi Grossman and George Fitzmaurice. 2010. ToolClips: An Investigation of Contextual Video Assistance for Functionality Understanding. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, Atlanta Georgia USA, 1515–1524. <https://doi.org/10.1145/1753326.1753552>
- [29] Transparent Statistics in Human-Computer Interaction Working Group. 2019. Transparent Statistics Guidelines. <https://doi.org/10.5281/zenodo.1186169>
- [30] Thomas Haigh. 2006. Remembering the Office of the Future: The Origins of Word Processing and Office Automation. *IEEE Annals of the History of Computing* 28, 4 (Oct. 2006), 6–31. <https://doi.org/10.1109/MAHC.2006.70>
- [31] Sarah Henderson. 2011. Document Duplication: How Users (Struggle to) Manage File Copies and Versions. *Proceedings of the American Society for Information Science and Technology* 48, 1 (2011), 1–10. <https://doi.org/10.1002/meet.2011.14504801013>
- [32] Josh Holinaty, Alec Jacobson, and Fanny Chevalier. 2021. Supporting Reference Imagery for Digital Drawing. In *2021 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*. IEEE, Montreal, BC, Canada, 2434–2442. <https://doi.org/10.1109/ICCVW54120.2021.00276>
- [33] Md Naimul Hoque, Tasfia Mashiat, Bhavya Ghai, Cecilia D. Shelton, Fanny Chevalier, Kari Kraus, and Niklas Elmqvist. 2024. The HaLLMark Effect: Supporting Provenance and Transparent Use of Large Language Models in Writing with Interactive Visualization. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–15. <https://doi.org/10.1145/3613904.3641895>
- [34] Dun-Ming Huang, Pol Van Rijn, Ilia Sucholutsky, Raja Marjeh, and Nori Jacoby. 2024. Characterizing Similarities and Divergences in Conversational Tones in Humans and LLMs by Sampling with People. <https://doi.org/10.48550/arXiv.2406.04278> arXiv:2406.04278 [cs]
- [35] Edwin L. Hutchins, James D. Hollan, and Donald A. Norman. 1985. Direct Manipulation Interfaces. *Human-Computer Interaction* 1, 4 (Dec. 1985), 311–338. https://doi.org/10.1207/s15327051hci0104_2
- [36] Takumi Ito, Tatsuki Kuribayashi, Masatoshi Hidaka, Jun Suzuki, and Kentaro Inui. 2020. Langsmith: An Interactive Academic Text Revision System. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Association for Computational Linguistics, Online, 216–226. <https://doi.org/10.18653/v1/2020.emnlp-demos.28>
- [37] Wallace Jackson. 2015. The Compositing of Digital Illustration: Using Layers. In *Digital Illustration Fundamentals: Vector, Raster, WaveForm, NewMedia with DICC, DAEF and ASNMF*, Wallace Jackson (Ed.). Apress, Berkeley, CA, 133–148. https://doi.org/10.1007/978-1-4842-1697-2_12
- [38] Di Jin, Zhijing Jin, Zhiting Hu, Olga Vechtomova, and Rada Mihalcea. 2021. Deep Learning for Text Style Transfer: A Survey. arXiv:2011.00416 [cs]
- [39] Chris Kim, Uta Hinrichs, Saif M. Mohammad, and Christopher Collins. 2020. Lexichrome: Text Construction and Lexical Discovery with Word-Color Associations Using Interactive Visualization. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference (DIS '20)*. Association for Computing Machinery, New York, NY, USA, 477–488. <https://doi.org/10.1145/3357236.3395503>
- [40] Yoonsu Kim, Jueon Lee, Seoyoung Kim, Jaehyuk Park, and Juho Kim. 2024. Understanding Users' Dissatisfaction with ChatGPT Responses: Types, Resolving Tactics, and the Effect of Knowledge Level. In *Proceedings of the 29th International Conference on Intelligent User Interfaces (IUI '24)*. Association for Computing Machinery, New York, NY, USA, 385–404. <https://doi.org/10.1145/3640543.3645148>
- [41] Cynthia L. King. 2012. Reverse Outlining: A Method for Effective Revision of Document Structure. *IEEE Transactions on Professional Communication* 55, 3 (Sept. 2012), 254–261. <https://doi.org/10.1109/TPC.2012.2207838>
- [42] Matthew G. Kirschenbaum. 2016. *Track Changes: A Literary History of Word Processing* (illustrated edition ed.). Belknap Press: An Imprint of Harvard University Press, Cambridge, Massachusetts.
- [43] Sofia Kitromili, James Jordan, and David E. Millard. 2019. What Is Hypertext Authoring?. In *Proceedings of the 30th ACM Conference on Hypertext and Social*

- Media (HT '19)*. Association for Computing Machinery, New York, NY, USA, 55–59. <https://doi.org/10.1145/3342220.3343653>
- [44] Susan Lin, Jeremy Warner, J.D. Zamfirescu-Pereira, Matthew G Lee, Sauhard Jain, Shanqing Cai, Piyawat Lertvittayakumjorn, Michael Xuelin Huang, Shumin Zhai, Bjoern Hartmann, and Can Liu. 2024. Rambler: Supporting Writing With Speech via LLM-Assisted Gist Manipulation. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–19. <https://doi.org/10.1145/3613904.3642217>
- [45] Wendy Mackay. 2000. Responding to Cognitive Overload : Co-adaptation between Users and Technology. *Intellectica* (2000), 177–193. <https://doi.org/10.3406/intel.2000.1597>
- [46] Cerstin Mahlow. 2023. Writing Tools: Looking Back to Look Ahead. arXiv:2303.17894 [cs]
- [47] Damien Masson, Sylvain Malacria, Géry Casiez, and Daniel Vogel. 2023. Statslator: Interactive Translation of NHST and Estimation Statistics Reporting Styles in Scientific Documents. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST '23)*. Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3586183.3606762>
- [48] Damien Masson, Sylvain Malacria, Géry Casiez, and Daniel Vogel. 2024. Direct-GPT: A Direct Manipulation Interface to Interact with Large Language Models. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–16. <https://doi.org/10.1145/3613904.3642462>
- [49] Damien Masson, Jo Vermeulen, George Fitzmaurice, and Justin Matejka. 2022. Supercharging Trial-and-Error for Learning Complex Software Applications. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3491102.3501895>
- [50] Jesse G. Meyer, Ryan J. Urbanowicz, Patrick C. N. Martin, Karen O'Connor, Ruowang Li, Pei-Chen Peng, Tiffani J. Bright, Nicholas Tatonetti, Kyoung Jae Won, Graciela Gonzalez-Hernandez, and Jason H. Moore. 2023. ChatGPT and Large Language Models in Academia: Opportunities and Challenges. *BioData Mining* 16, 1 (July 2023), 20. <https://doi.org/10.1186/s13040-023-00339-9>
- [51] Dennis C. Neale and John M. Carroll. 1997. Chapter 20 - The Role of Metaphors in User Interface Design. In *Handbook of Human-Computer Interaction (Second Edition)*, Marting G. Helander, Thomas K. Landauer, and Prasad V. Prabhur (Eds.). North-Holland, Amsterdam, 441–462. <https://doi.org/10.1016/B978-044481862-1.50086-8>
- [52] nextui-org 2021. *NextUI* nextui-org. <https://nextui.org/>
- [53] Jakob Nielsen. 1994. *Usability Engineering* (1st edition ed.). Morgan Kaufmann, Boston.
- [54] Ayano Okoso, Keisuke Otaki, Satoshi Koide, and Yukino Baba. 2024. Impact of Tone-Aware Explanations in Recommender Systems. <https://doi.org/10.48550/arXiv.2405.05061> arXiv:2405.05061 [cs]
- [55] OpenAI. 2023. OpenAI Node API Library. OpenAI.
- [56] Alice Oshima and Ann Hogue. 2005. *Writing Academic English* (4th edition ed.). Pearson Education ESL, White Plains, NY.
- [57] P. G. Polson, E. Muncher, and G. Engelbeck. 1986. A Test of a Common Elements Theory of Transfer. *ACM SIGCHI Bulletin* 17, 4 (April 1986), 78–83. <https://doi.org/10.1145/22339.22352>
- [58] Jenny Preece, Yvonne Rogers, and Helen Sharp. 2002. *Interaction Design: Beyond Human-Computer Interaction* (1st edition ed.). Wiley, New York, NY.
- [59] Jef Raskin. 2000. *The Humane Interface: New Directions for Designing Interactive Systems* (1st edition ed.). Addison-Wesley Professional, Reading, Mass.
- [60] Meta 2013. *React – A JavaScript Library for Building User Interfaces*. Meta. <https://reactjs.org/>
- [61] Miguel A. Renom, Baptiste Caramiaux, and Michel Beaudouin-Lafon. 2023. Interaction Knowledge: Understanding the 'Mechanics' of Digital Tools. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, Hamburg Germany, 1–14. <https://doi.org/10.1145/3544548.3581246>
- [62] Mohi Reza, Nathan Laundry, Ilya Musabirov, Peter Dushniku, Zhi Yuan "Michael" Yu, Kashish Mittal, Toví Grossman, Michael Liut, Anastasia Kuzminykh, and Joseph Jay Williams. 2024. ABScribe: Rapid Exploration & Organization of Multiple Writing Variations in Human-AI Co-Writing Tasks Using Large Language Models. <https://doi.org/10.1145/3613904.3641899> arXiv:2310.00117 [cs]
- [63] Yvonne Rogers, Helen Sharp, and Jennifer Preece. 2023. *Interaction Design: Beyond Human-Computer Interaction* (6th edition ed.). Wiley, Hoboken.
- [64] Mike Rose. 1980. Rigid Rules, Inflexible Plans, and the Stifling of Language: A Cognitivist Analysis of Writer's Block. *College Composition and Communication* 31, 4 (1980), 389–401. <https://doi.org/10.2307/356589> jstor:356589
- [65] Literature & Latte Ltd 2005. *Scrivener*. Literature & Latte Ltd. <https://www.literatureandlatte.com/scrivener>
- [66] Kevin Sheppard, Stanislav Khrapov, Gábor Lipták, mikedelalima, Rob Capellini, alejandro cermeno, Hugle, esvhd, Snyk bot, Alex Fortin, JPN, Matt Judell, Ryan Russell, Weiliang Li, 645775992, Austin Adams, jbrockmendel, LGTM Migrator, M. Rabba, Michael E. Rose, Nikolay Tretyak, Tom Rochette, UNO Leo, Xavier RENE-CORAIL, Xin Du, Joren, and Burak Çelik. 2023. *bashtage/arch: Release 6.1.0*. <https://doi.org/10.5281/zenodo.7975104>
- [67] Hirohito Shibata and Koichi Hori. 2008. Cognitive Support for the Organization of Writing. *New Generation Computing* 26, 2 (Feb. 2008), 97–124. <https://doi.org/10.1007/s00354-008-0037-9>
- [68] Shneiderman. 1983. Direct Manipulation: A Step Beyond Programming Languages. *Computer* 16, 8 (Aug. 1983), 57–69. <https://doi.org/10.1109/MC.1983.1654471>
- [69] Ben Shneiderman, Gerhard Fischer, Mary Czerwinski, Mitch Resnick, Brad Myers, Linda Candy, Ernest Edmonds, Mike Eisenberg, Elisa Giaccardi, Tom Hewett, Pamela Jennings, Bill Kules, Kumiyo Nakakoji, Jay Nunamaker, Randy Pausch, Ted Selker, Elisabeth Sylvan, and Michael Terry. 2006. Creativity Support Tools: Report From a U.S. National Science Foundation Sponsored Workshop. *International Journal of Human-Computer Interaction* 20, 2 (May 2006), 61–77. https://doi.org/10.1207/s15327590ijhc2002_1
- [70] Maria Shugrina, Wenjia Zhang, Fanny Chevalier, Sanja Fidler, and Karan Singh. 2019. Color Builder: A Direct Manipulation Interface for Versatile Color Theme Authoring. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3290605.3300686>
- [71] Slate.js 2016. *Slate.js*. Slate.js. <https://www.slatejs.org>
- [72] Sangho Suh, Meng Chen, Bryan Min, Toby Jia-Jun Li, and Haijun Xia. 2024. Luminate: Structured Generation and Exploration of Design Space with Large Language Models for Human-AI Co-Creation. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–26. <https://doi.org/10.1145/3613904.3642400>
- [73] Theophanis Tsandilas, Magdalini Grammatikou, and Stéphane Huot. 2015. BricoSketch: Mixing Paper and Computer Drawing Tools in Professional Illustration. In *Proceedings of the 2015 International Conference on Interactive Tabletops & Surfaces (ITS '15)*. Association for Computing Machinery, New York, NY, USA, 127–136. <https://doi.org/10.1145/2817721.2817729>
- [74] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, António H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. 2020. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* 17 (2020), 261–272. <https://doi.org/10.1038/s41592-019-0686-2>
- [75] Amelia Wattenberger. 2024. *GitHub Next | Code Brushes*. <https://next.github.com/projects/code-brushes>.
- [76] Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, Myra Cheng, Mia Glaese, Borja Balle, Atoosa Kasirzadeh, Zac Kenton, Sasha Brown, Will Hawkins, Tom Stepleton, Courtney Biles, Abeba Birhane, Julia Haas, Laura Rimel, Lisa Anne Hendricks, William Isaac, Sean Legassick, Geoffrey Irving, and Iason Gabriel. 2021. Ethical and Social Risks of Harm from Language Models. <https://doi.org/10.48550/arXiv.2112.04359> arXiv:2112.04359 [cs]
- [77] Brian Wood. 2022. *Adobe Illustrator Classroom in a Book* (1st edition ed.). Adobe Press.
- [78] Haijun Xia, Bruno Araujo, Toví Grossman, and Daniel Wigdor. 2016. Object-Oriented Drawing. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, San Jose California USA, 4610–4621. <https://doi.org/10.1145/2858036.2858075>
- [79] Polle T. Zellweger, Susan Harkness Regli, Jock D. Mackinlay, and Bay-Wei Chang. 2000. The Impact of Fluid Documents on Reading and Browsing: An Observational Study. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, The Hague The Netherlands, 249–256. <https://doi.org/10.1145/332040.332440>
- [80] Zheng Zhang, Jie Gao, Ranjodh Singh Dhaliwal, and Toby Jia-Jun Li. 2023. VISAR: A Human-AI Argumentative Writing Assistant with Visual Programming and Rapid Draft Prototyping. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. ACM, San Francisco CA USA, 1–30. <https://doi.org/10.1145/3586183.3606800>
- [81] Yaqian Zhu and John Kolassa. 2018. Assessing and Comparing the Accuracy of Various Bootstrap Methods. *Communications in Statistics - Simulation and Computation* 47, 8 (Sept. 2018), 2436–2453. <https://doi.org/10.1080/03610918.2017.1348516>
- [82] William Zinsser. 2016. *On Writing Well: The Classic Guide to Writing Nonfiction* (30th anniversary ed. edition ed.). Harper Perennial, New York London Toronto Sydney New Delhi Auckland.