Kandinsky Meets Social Conversation: Towards Abstract Art-Inspired Visualization Abstraction for Mobile Devices

Castillo Clarence Fitzerald Gumtang, Chun Leong Tham, Sourav S. Bhowmick
Nanyang Technological University (NTU), Singapore
clarence001@e.ntu.edu.sg, ctham005@e.ntu.edu.sg, assourav@ntu.edu.sg

Abstract

Social media platforms enable users and organizations to share content–such as text, images, and videos–collectively referred to as anchor posts. These posts often attract hundreds of comments, forming rich social conversations. Traditionally, these comment threads are displayed in a hierarchical list format, which users must manually browse to understand the discussion dynamics. However, this process can be tedious and inefficient, discouraging users from conducting multi-faceted exploration of the conversation landscape. In this paper, we introduce a novel visualization abstraction called *Several Concentric Circles (SCC)*, which transforms textual social discussions to a digital "artwork" to provide an interactive, mobile-friendly, and visually appealing alternative to conventional list-based abstraction. SCC addresses the limitations of traditional formats by integrating principles from two seemingly distinct domains–abstract art and computer science–to foster both aesthetic appeal and user engagement. To validate our approach, we developed a prototype system named Kandinsky Mobile, which brings our visualization framework to life and supports effective exploration of social discussions.

1 Introduction

Social media platforms (*e.g., Facebook*, *YouTube*, *X*) enable individuals and organizations to upload content such as posts, images, or videos–collectively referred to as *anchor posts*. Other users can engage with these posts by reading and commenting on them, resulting in social conversations (*i.e.,* discussions). In essence, anchor posts are essentially social media posts with accompanied comments, facilitating discussion and interaction. Posts can feature various elements, including visuals, text, and interactive components, to capture user interest and build community engagement. Comments typically appear in hierarchical list-based format and serve as textual or multimedia responses that allow users to share opinions, pose questions, and take part in the ongoing conversation. With the rapid rise of mobile technology, these interactions increasingly occur via mobile devices. For instance, mobile phones are often the primary means by which users access *X*[1].

A large number of online users read comments associated with an anchor post for various reasons such as to learn about opinions of others, to be entertained or amused, to get information, to compare one's view on a topic with others, etc [7]. Unfortunately, this can be cognitively challenging when one has to peruse a large number of comments. To mitigate this cognitive burden, most social media platforms automatically display comments deemed *"top"*, *"most relevant"*, or *"most recent"*, as determined by proprietary machine learning (ML) algorithms. These algorithms are typically optimized to maximize user engagement and platform profitability [35]. It is worth noting that in recent times, such ML-driven

[1] https://blog.X.com/en_us/a/2013/new-compete-study-primary-mobile-users-on-twitter

exposition models in social media sites (*e.g., Facebook*) have led to *"proliferation of extremism, bullying, hate speech, disinformation, conspiracy theory, and rhetorical violence"* [35].

1.1 Motivation

Despite the prevalence of the list-based visualization abstraction of social media discussions with the goal to optimize engagement and profit, it suffers from several key limitations as follows.

- First, reading beyond the top or most relevant/recent comments can be tedious and time-consuming, particularly when an anchor post has a large volume of responses. In many ways, this process resembles navigating a lengthy hierarchical text document. This is especially aggravated in a mobile device due to its small screen size and unique interaction style [10, 26, 29]. In particular, none of the social media sites allow a user to get a bird's-eye view of the conversation landscape, highlighting discussion threads that have garnered considerable attention and interactions between various comments and users. Such view can enable users to choose *efficiently any* conversation that they *wish* to peruse instead of steered by ML-based algorithms.

- Second, users are generally unable to perform interactive search within the discussion threads, making it difficult to locate specific topics, keywords, or user contributions across potentially vast comment sections. Additionally, there is limited support for temporal navigation, what we refer to as "time traveling" through the discussion space, where users could explore how conversations evolved over time or revisit specific moments in the discussion's history. This lack of temporal and semantic interactivity significantly hinders deeper exploration, analysis, and understanding of social discourse.

- Third, while aesthetics is a well-established criterion for usability [13] and a key component of effective human-data interaction frameworks [50], the conventional list-based presentation of social media comments tends to lack aesthetic appeal as it prioritizes individual comment readability over aesthetics. This linear, text-heavy format can appear cluttered, monotonous, and visually unengaging, which may detract from the overall user experience and hinder meaningful engagement with the content.

- Last but not the least, ML-based exposition of *relevant* comments, while designed to surface the most engaging or relevant content, can inadvertently lead users into "filter bubbles" [35]. These algorithms tend to prioritize comments that align with a user's prior interactions, preferences, or inferred interests, thereby limiting exposure to diverse viewpoints within the broader discussion landscape. As a result, users may be systematically steered away from diverse opinions or novel perspectives, undermining the potential for open dialogue, critical engagement, and the healthy exchange of ideas.

To address the aforementioned limitations, at a high-level we aim to *transform* the hierarchical textual comments linked to an anchor post into an aesthetically pleasing, interactive "artwork". This concept has gained momentum recently with the emergence of text-to-image generation systems [4, 6], which are capable of producing digital images and artworks from natural language prompts. To this end, we introduce a novel *abstract art-inspired visualization abstraction* for representing social discussions centered around an anchor post. It enables us to achieve the following set of key visualization goals while laying the groundwork for addressing the aforementioned limitations.

1. *Canvas Construction:* Aesthetically pleasing, complete, and mobile-friendly bird's eye view of the social discussions landscape.

2. *On-demand Conversation Visualization:* On-demand, mobile-friendly visualization and exploration of comments and discussion threads.

3. *Canvas Search:* Interactive keyword and comment similarity search on the discussions and visualization of search results.

4. *Time Travel:* Time travel of the discussions landscape.

Specifically, Goals 1 and 2 address the first and third limitations, while Goals 3 and 4 target the second. Together, all four are intended to establish a foundation for tackling the final limitation.

1.2 Overview

Devising visualization abstraction that realize the aforementioned goals is challenging. While challenges such as on-demand exploration, interactive search, mobile-friendly visualization, and temporal navigation fall squarely within the domain of computer science, the dimension of aesthetics is traditionally rooted in the arts. Consequently, we seek inspirations from the abstract arts of the famous Russian painter Wassily Kandinsky (1866-1944) to build a novel visualization abstraction called *Several Concentric Circles* (SCC) that lays the foundation for transforming the textual comment space of an anchor post to a digital, interactive artwork. The data model of SCC consists of a set of colorful *circles* and *shape buckets* that contain these circles. Each circle represents a comment in a social discussion. The social conversation landscape with respect to an anchor post is then *transformed* to an aesthetically pleasing, interactive collection of colorful circles and concentric circles on a mobile display canvas inspired by the artworks of Kandinsky using a set of *operations*. In particular, more users are involved in a discussion, the larger and more colorful the concentric circles are on the canvas. Specifically, these operations facilitate realization of all the aforementioned visualization goals. We also introduce a prototype implementation on *YouTube*, named Kandinsky-Mobile[2] , which realizes the SCC abstraction to demonstrate the proposed visualization goals. It is worth noting that while text-to-image generation systems have gained increasing popularity with the rise of generative AI, Kandinsky-Mobile [22] and its desktop variant [38] predate these systems. In summary, this paper makes the following contributions.

- We propose a novel visualization abstraction, *Several Concentric Circles (SCC)*, inspired by abstract arts of Wassily Kandinsky, to "paint" the landscape of social discussions surrounding an anchor post. To the best of our knowledge, this is the first work to bridge the domains of abstract art and data visualization in the context of visualizing social media conversations.

- To systematically realize the aforementioned visualization goals, we introduce a set of novel *SCC operations*–a collection of core interaction and rendering operations grounded in the design principles of our SCC abstraction. These operations serve as the foundation for achieving the aforementioned goals. Building on these operations, we develop Kandinsky Mobile, a prototype system that demonstrates the practical viability of our approach on mobile devices.

- Our initial experimental evaluation and user study highlight the effectiveness and superiority of SCC compared to traditional list-based visualization abstraction in achieving the aforementioned goals.

The rest of the paper is organized as follows. Section 2 reviews related work. Section 3 presents a brief background on Kandinsky's artworks and inspiration behind our SCC abstraction and Kandinsky

[2] Kandinsky-Mobile and its desktop variant are demonstrated in [22] and [38], respectively.

Mobile. Sections 4 and 5 describe the architecture of Kandinsky Mobile and the SCC abstraction that underpins it. We describe how the aforementioned visualization goals are realized using SCC in Section 6. We report our preliminary performance study in Section 7. The final section concludes the paper with reflections on the proposed visualization paradigm and outlines open challenges and future directions in this interdisciplinary area.

2 Related Work

Conversation visualization. Most germane to our research are efforts on visualizing social comments and conversations. Research on this arena dates back to the 1990s and primarily focuses on identifying the temporal or structural relationship between conversations.

Donath *et al.* [14] proposed *Chat Circles*, a graphical interface for visualizing synchronous conversations. The interface represents each participation by a colored circle on the screen in which his or her posts appear. The interface *Conversation Landscape* was also developed to visualize the conversational archive of *Chat Circles*. It provides a 2-dimensional visualization of a chat group where the x-axis represents participants identified by color and the y-axis represents timeline of postings by participants. Each post is shown as a horizontal line, with the length reflecting the message's length. This design provides a bird's-eye view of the activeness of participants in the discussion over time and the interaction patterns of the conversation, making dominant users and surge of discussions recognizable. The limitation of this design is that, for discussion among a large number of participants but with low contribution rate, it may render a long horizontal axis but sparse lines.

Conversation Thumbnails [51] visualizes the thread structure of conversations using tree-visualization techniques. Engdahl *et al.* [17] used *Squarified Treemap Layout* to visualize threaded discussion forums on devices with small screen size. The technique uses squarified treemaps to display the contents of discussion forums. Each thread from discussion forums is rendered as a colored rectangle in a treemap. The size of the rectangle is either proportional to the number of articles in the thread, or reflects the relevance of the query for that thread when searching is performed. The user is first presented an overview of the top level threads for each forum. As the stylus is moved across the screen, the details of the selected thread is displayed in the popup. This approach makes use of 100% of the limited screen space and provides an intuitive overview of the discussion landscape, allowing the user to easily compare popularity among threads. The most (resp. less) discussed thread is placed in the upper left (resp. bottom right) corner of the map. However, the ability to retrieve detailed information and analyze temporal trends is missing.

PeopleGarden [52] creates a garden of flowers to visualize users' interactions over time. Each flower represents a user and each petal is a post by him/her. Each user is portrayed as a flower, whose petals represent the user's posts, and height reflecting the amount of time the user has been in the discussion. A flower "blossoms" as a user becomes more engaged in a discussion. Together, users' interactions in the social environment form a garden. This design achieves compact information representation about participants involved in online discussions, but fails to provide details related to the discussion. Furthermore, flowers might overlap and become difficult to analyze once the more participants are involved.

ToPIN [48] is a system for analyzing and visualizing comments on online educational videos. It provides a qualitative perspective by examining the specifics of individual comments, including their type, timestamp, and how they cluster and evolve over time. Additionally, it incorporates *ThemeRiver* [23] to offer a complementary quantitative view, highlighting trends and shifts within the collection of time-stamped comments.

ConVis [27] visualizes blog conversations through a thread overview that reveals their structure and sentiment, alongside a faceted overview that highlights key topics and contributors. More recently, [49] used sentiment analysis on *Yahoo! News* article comments about Covid19 testing, classifying them as negative, positive, or neutral, and leverage a bubble chart for visualizing public sentiment and its monthly changes based on response volume and approval rates. *EmoVis* [28] is a framework for visualizing sentiment shifts in social network comment sections as a way to track public discourse over time. Using a pre-trained model, it predicts emotional scores from user comments and maps them to various sentiment categories. However, this framework does not propose any novel visualization abstraction. These efforts are complimentary to our work as we do not focus on visualizing sentiments.

None of the aforementioned visualization schemes for conversation draw inspiration from abstract art. Moreover, they lack a systematic implementation of visualization abstractions based on a defined data model and a corresponding set of operations. In contrast, Kandinsky Mobile utilizes the proposed SCC abstraction to design and build an interactive visualization framework that represents the landscape of social discussions on mobile devices.

Text-to-artwork generation. With the emergence of *Generative Artificial Intelligence (GenAI)*, several large-scale text-to-art models and systems have been introduced [4, 6]. A shared characteristic among these models is their capability to generate digital artwork based on user-written prompts in natural language. In a similar vein, Kandinsky Mobile seeks to convert textual user comments, rather than explicit prompts, into visually appealing artwork, while avoiding reliance on large generative models. Additionally, unlike Kandinsky Mobile, these systems are not tailored to a specific artistic style, such as that of Kandinsky.

Recently, Zhou *et al.* [55] developed a text-to-image dataset centered on the works of Wassily Kandinsky and introduced a generative method that allows users to automatically produce personalized artwork in Kandinsky's style, based on their preferences–removing the need for explicit prompts. In contrast to SCC, their approach does not rely on any data model for creating artwork in Kandinsky's style. The objectives are also different. While they focus on generating personalized digital artworks in the style of Kandinsky, our aim is to create Kandinsky-inspired visual abstractions for social discussions.

Complex circle-based network data visualization. The use of complex circular geometries is well-established within the visualization community, particularly for illustrating various properties and structures of networks [9, 42, 54]. For instance, *hybrid layout* techniques, as described in [42], combine multiple layout strategies to effectively capture and display the topological features of networks. In a related approach, the work presented in [9] introduces composite circular encodings, where *Treemap visualizations*, typically used for hierarchical data, are integrated within circular frameworks to represent different network layers within a unified visual metaphor. Similarly, the *Graph Thumbnails* technique [54] leverages circle packing to provide an overview visualization of hierarchical network structures. In the domain of property graph visualization, the concept of *composite* circles (*i.e.,* primary circles surrounded by smaller satellite circles) has been recently utilized to visually encode both property graph queries and schema structures [11, 40], offering an intuitive and spatially compact representation of complex property graph semantics. These efforts are primarily driven by structural, semantic, or functional considerations and are not informed by any particular tradition of abstract artistic expression. The SCC abstraction draws deliberate inspiration from abstract art. Specifically, it employs vibrant, colorful concentric circles as a central visual motif to aesthetically represent the dynamic conversations among social media users in relation to a given anchor post.

User interface design for mobile devices. The need to find effective user interfaces for mobile devices is critical. In order to save space on small screens, Noirhomme-Fraiture *et al.* [43] suggested to use concise and precise information. Giller *et al.* [20] observed that users performed significantly better when they could scroll, instead of tapping on widget elements tabs for page-to-page navigation, and that

Figure 1: **[Best viewed in color]** Artworks of Wassily Kandinsky: (Left) *Squares with Concentric Circles* (1913); (Middle) *Several Circles* (1926) [2]; (Right) *Heavy Circles* (1927) [3].

scrolling vertically rather than horizontally led to better user experience. Vertical scrolling should be used when displaying the long list of comments, such that the number of taps can be minimized [46]. To overcome the limited interaction with a finger touch (*i.e.,* the thumb), the application must be fat-finger friendly [5, 10, 45]. Due to the limited motion available to the thumb, Hoober [24, 25] observed based on collected data that users prefer to view and touch the center of the screen. People prefer reading content at the center of the screen and are better at tapping at the center of the screen. Kandinsky Mobile embodies these key insights through the implementation of the SCC operations.

3 Inspirations Behind *Kandinsky Mobile*

In this section, we present relevant abstract artworks of Wassily Kandinsky that serve as inspiration to our SCC abstraction. In particular, we articulate how his artistic philosophy in these artworks informed the design of Kandinsky Mobile. In the subsequent sections, we shall elaborate on the realization of this design philosophy.

3.1 Wassily Kandinsky: Father of Abstract Art

As the "father" of abstract art, Wassily Kandinsky (1866-1944) is known for his lyrical style and innovative theories on non-representational art[3]. As opposed to traditional realism that focused on physical matter and materialism, he exploited the evocative interrelation between *color* and *form* to create an aesthetic experience. A recent study reported that Kandinsky's art has impacted aesthetic experience emphasizing that modern art is not meaningless and has value because of its artistic expression [12].

Color. Kandinsky's color theory is built around three primary colors, *Red*, *Yellow*, and *Blue*, and three primary tones, *Black*, *White*, and *Grey* [32]. By combining the primary colors, he derives three secondary colors, *Orange*, *Green*, and *Violet*. He classifies these six chromatic colors based on perceived temperature: red, yellow, and orange are categorized as *warm*, while green, blue, and violet are considered *cold*. Tone in Kandinsky's framework refers to the lightness or darkness of a composition, determined by the ratio of black to white. Each color is associated with a distinct spiritual emotion and exerts an influence on surrounding colors [32], similar to how elements in a visual or perceptual system interact and modify one another. This anticipates modern theories in computational aesthetics and HCI, where color is used not just for visibility but also for emotional engagement and user experience design.

[3]Representational art is designed to represent real life, and non-representational art is the opposite.

Form. Geometrical elements took on increasing importance in Kandinsky's painting – particularly the circle. He proposed a structured theory of visual tension, defining graphical elements in terms of their shape, spatial position, and orientation. He begins with the primitive unit, the *Point*, and derives two higher-order constructs: the *Line* and the *Plane* [33]. Lines are categorized based on the type of tension they exhibit, resulting in three forms: *Straight*, *Curved*, and *Angular* [33]. Similarly, planes are subject to tension and can be represented as *Triangles*, *Squares*, or *Circles* [33]. These seven geometric primitives serve as the core components in Kandinsky's visual compositions, akin to how basic shapes form the building blocks in computational visual systems.

Kandinsky asserts that the overall composition is shaped by the dynamic interrelationships among all visual elements. It is defined not only by the placement of individual elements but also by the connections they form with one another and the structure of grouped elements within the visual space.

3.2 Circles in Kandinsky's Artworks

The aforementioned artistic philosophy of Kandinsky can be experienced in his works such as *Circles in a Circle* (1923) [1], *Several Circles* (1926) [2], and *Heavy Circles* (1927) [3] (Figure 1). We focus on these as they serve as inspiration to the design of SCC and Kandinsky Mobile.

Squares with Concentric Circles (1913). *Squares with Concentric Circles* (Figure 1 (left)) is a key work by Kandinsky that exemplifies his use of abstraction, color theory, and spiritual symbolism. The painting features concentric circles within squares, where the square acts as a static container, providing contrast to the dynamic, fluid nature of the circles. This creates a tension between the infinite (circle) and the finite (square), a recurring theme in Kandinsky's work that explores the balance between opposing forces, such as the spiritual vs. the material and freedom vs. constraint.

The concentric circles are layered with different colors, which not only create visual movement but also evoke emotional depth. While the circles vary in color and size, they maintain harmony within the square, reflecting Kandinsky's belief in balance and the significance of geometric forms in expressing emotional and spiritual resonance.

Several Circles (1926). Kandinsky regarded the circle as the most "spiritual" of all geometric forms. In *Several Circles* (Figure 1 (middle)), he isolates this shape to investigate its expressive possibilities. Unlike his earlier works, which often included a variety of shapes and intersecting lines, this work is dedicated solely to only one form, the circle. This could focus instead on colour, mass and the relative positions of the circles, minimizing visual noise and emphasizing their symbolic meaning. The circles differ in size, color, transparency, and placement, generating a visually intricate yet harmonious composition. This approach reflects Kandinsky's belief that effective composition arises not from symmetry, but from the dynamic interplay between forms. Each circle interacts with the others through its color and placement, illustrating his idea that elements in a composition influence one another.

Kandinsky also applies his color theory to these forms, using contrasts of temperature (warm vs. cool) and tone (light vs. dark) to evoke emotional intensity. Color affects not only the individual circle's feel but also its relationship to the rest of the composition. The largest circle, surrounded by a glowing halo, commands attention and becomes a focal point. The careful positioning of each circle reflects Kandinsky's theory of pictorial tension, where shapes exert visual "forces" on one another based on their size, orientation, and location. This interaction generates a sense of dynamic equilibrium–one that feels alive despite the static nature of the shapes.

Heavy Circles (1927). In *Heavy Circles* (Figure 1 (right)), Kandinsky builds on his exploration of geometric abstraction, emphasizing visual weight, density, and contrast. Similar to *Several Circles*, the composition is centered around two large circles, with a series of smaller circles placed in the corners of the artwork. The painting introduces variations in line, color, and complexity, differentiating it from his

Table 1: Relationship between Kandinsky's artworks and the features of Kandinsky Mobile.

ID	*Artwork*	*Feature in Kandinsky's artworks*	*Feature in Kandinsky Mobile*	*Semantics*	*SCC Operations*
K1	Several Circles, Heavy Circles	Circles are the only geometric forms to minimize visual noise.	Circles are the only geometric shape in SCC.	Each circle represents a comment.	paint
K2	Squares with Concentric Circles	A collection of colorful concentric circles layered with different colors to create visual movement and emotional depth.	Nucleus concentric circles layered with different colors.	Represents replies to a comment.	paint
K3	Several Circles, Heavy Circles	Circles of different size and color may overlap on the circumference.	Peripheral concentric circles.	Represents replies to a reply (depth greater than or equal to 2 in a conversation network).	paint
K4	Squares with Concentric Circles, Several Circles, Heavy Circles	Usage of different colors, contrast of temperature and tone.	Color coding of circles using different temperature and tone.	Different users in a conversation.	paint
K5	Several Circles, Heavy Circles	Two larger circles are the focal points.	Larger concentric circles are focal points.	Represents conversations that have garnered a lot of interactions.	paint, balance
K6	Several Circles, Heavy Circles	Circles are strategically positioned to minimize visual clutter and highlight interplay between forms.	Circles are positioned strategically to reduce cluttering and showing interactions in a conversation.	Represents the interactions in the conversation space.	balance, update
K7	Squares with Concentric Circles	Squares encapsulate the concentric circles.	Squares are used to encapsulate circles that satisfy search results.	Represents results of interactive search.	update
K8	Several Circles, Heavy Circles	Dark background of the canvas.	Dark background of *Kandinsky Canvas*.	Choice of background color for Kandinsky Mobile.	paint
K9	Heavy Circles	Color intensity is used to show the interplay of heavier and lighter circles.	Color intensity is lowered in interactive search and time travel.	Lower intensity circles represent conversations that are not part of search results or not within a specific time window.	paint, update

earlier work *Several Circles*. The circles in *Heavy Circles* have a "heavier" appearance compared to the lighter, airier forms in *Several Circles*. Kandinsky manipulates color saturation, size, and placement to create a sense of mass, making the shapes feel anchored and impactful, both visually and psychologically. His theory of visual tension is evident in the balance between heavier and lighter circles. The "heaviness" of some areas is counteracted by lighter, more dispersed forms, creating a dynamic equilibrium that guides the viewer's eye. The interplay of color, size, and opacity also adds emotional depth, with darker tones evoking seriousness and brighter accents adding intensity.

3.3 Kandinsky's Circles to Kandinsky Mobile

Our system draws inspiration from these colorful, concentric, and overlapping circles in the artworks of Kandinsky. Table 1 gives an overview of the "mapping" between the three aforementioned artworks of Kandinsky and corresponding features in SCC and Kandinsky Mobile. The key features of these artworks that inspire the design of Kandinsky Mobile are as follows:

- Both *Several Circles* and *Heavy Circles* center around a single geometric element–the circle–to minimize visual noise. Likewise, the SCC abstraction exploits only circles to visualize the social conversation landscape (**K1**, **K3**).

- Variations in size, color, placement, and transparency of circles create a dynamic yet cohesive visual interaction in the artworks. The circles may be concentric or overlapping. Similarly, Kandinsky Mobile utilizes these same visual attributes–size, color, and placement–to depict the structure and semantics of social conversations (**K2-K4**, **K6**).
- In *Several Circles* and *Heavy Circles*, two prominent circles dominate the composition, serving as visual focal points. In parallel, Kandinsky Mobile positions large circles to highlight conversations that have attracted substantial engagement (**K5**).
- Kandinsky's approach to color resonates with contemporary digital color models. His use of warm versus cool tones and varying saturation not only generates visual rhythm but also communicates psychological states. Similarly, Kandinsky Mobile uses dark "canvas" and applies warm and cool colors along with saturation levels to encode information about the conversational landscape. For instance, color saturation is used for conveying time travel (**K8**, **K9**).
- The structured arrangement of concentric circles within squares in *Squares with Concentric Circles* informs our design choice to highlight results of interactive search in static square containers (**K7**).

4 The Several Concentric Circles (SCC) Abstraction

In this section, we present an *abstract art-inspired abstraction* of the social discussion landscape called *Several Concentric Circles (SCC)* whose design is inspired by the aforementioned artworks of Kandinsky. We begin with the architecture of Kandinsky Mobile system that realizes the SCC abstraction.

4.1 Architecture

Kandinsky Mobile is intended for users with normal vision who are experienced with mobile devices and social media conversations. Figure 2 depicts the architecture. A comment associated with a social post (*i.e., anchor post*) comprises the content, the author information (identifier and name), the date/time of comment, number of likes it has garnered (if any), and the *reply id* in which the comment is replied to. A set of such comments among individuals engenders social discussions. The goal of Kandinsky Mobile is to facilitate the four visualization goals introduced in Section 1 w.r.t the discussions associated with an anchor post.

Given a user-selected anchor post, all comments associated with it and their attributes are retrieved by the *Retrieval Service* via the social media data API. These comments and aforementioned attributes are stored in a key-value store by the *Persistence Layer* and indexes are created for efficient access to them. Furthermore, for each comment, topics are identified using Latent Dirichlet Allocation (LDA). The *Operator Layer* is responsible for the implementation of the *abstract art-inspired visualization abstraction* of Kandinsky Mobile. Specifically, it realizes the *map*, *paint*, *balance*, *retrieve*, *lookup*, and *update* operations which we shall describe in the sequel. The *Service Layer* exploits these operations to support the four goals of visualization, namely, *canvas construction*, *on-demand conversation visualization*, *canvas search*, and *time travel* (Recall from Section 1.1). A user interacts with these services through the *Presentation Layer*.

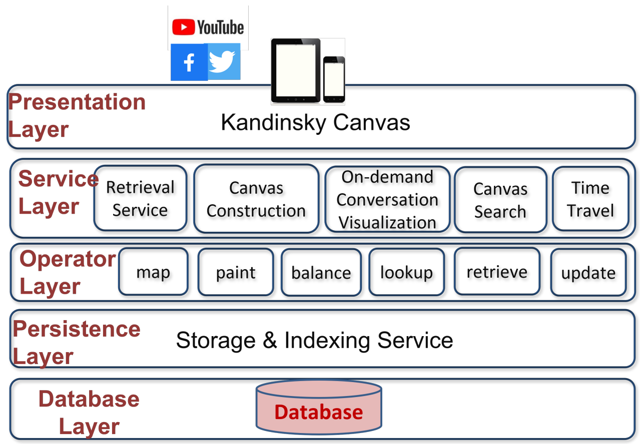

Figure 2: Architecture of Kandinsky Mobile.

4.2 The Data Model

We now present the SCC abstraction on which Kandinsky Mobile is built. Specifically, it aims to address the limitations of list-based abstractions highlighted earlier. At first glance, it may seem that we can address these limitations by switching to a directed graph-based data model where nodes represent comments and edges represent interactions between them (*e.g.,* replies, mentions, retweets). Although such purely computing-driven abstraction is reasonable to visualize social conversation, it becomes cognitively challenging to an end user when the graph size becomes large especially in small-screen mobile devices. In particular, large and dense graphs overload the human perception and cognitive systems, resulting in poor performance of relatively complex tasks such as identifying relationship in graphs [30, 53]. The study by [53] revealed that a graph with more than 100 edges can induce significant cognitive load on the end users. Unfortunately, many online discussion threads of anchor posts involve hundreds to thousands of comments.

The common limitation of list- or graph-based visualization abstraction is their ineffectiveness in exposing aesthetically pleasing bird's eye view of the discussion landscape that may galvanize end users to explore *any* part of the *entire* space on-demand. Hence, we depart from such traditional abstraction schemes and *explore the possibility of drawing inspirations from art to design a data model and operations on social discussions.* The rationale for pursuing this direction is aesthetics and engagement are closely associated with music and art. Integrating them with computational paradigms pave the way to address these limitations effectively.

At first glance, it may seem that the marriage of two disciplines that are poles apart, abstract art and computing, is quixotic. Certainly, as remarked in Section 2, this issue is hardly explored in the context of social discussion visualization. As an initial proof that such endeavour is indeed within the realm of possibility, we draw inspirations from the aesthetic art forms of Kandinsky (Section 3) for designing an abstraction that potentially leads to aesthetically pleasing and engaging visualization of a social discussion landscape.

Several Concentric Circles Data Model (SCC). Every abstraction in computer science consists of a data model and a way of manipulating the data. Social discussions typically adopt a hierarchical structure with the *original comment* (*i.e.,* comment which directly responds to an anchor post) as the root. We refer to it as *conversation network*. Figure 3(a) is an example of a conversation network rooted at Comment 1 where ids to comments are assigned according to their creation time. The SCC data model for visualizing the conversation networks of an anchor post (*i.e.,* discussion landscape) consists of the following:

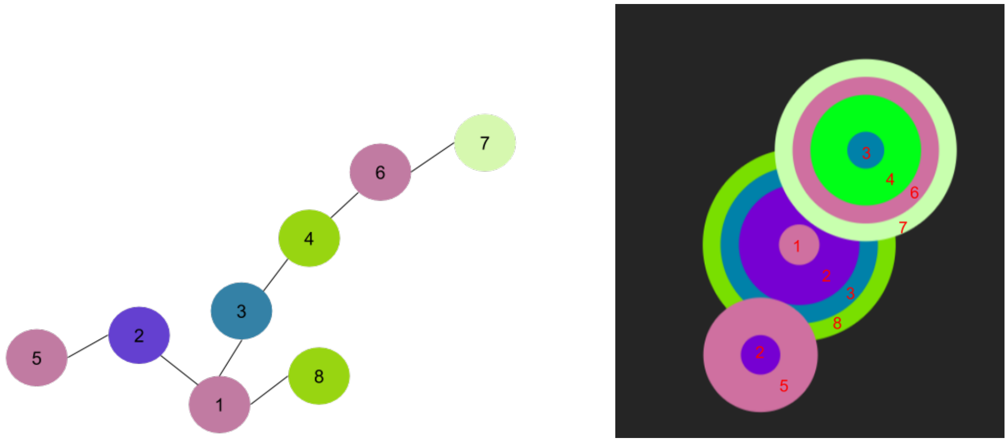

Figure 3: **[Best viewed in color]** A conversation network and its SCC abstraction.

1. A collection of colored shapes C. Each shape $c \in C$ represents a comment in a conversation network.

2. A set of *shape buckets* S (bucket for brevity) where $|S| \leq |C|$ in a 2D display space with dimension D.

3. A *map* function map from C to S. Each shape bucket s contains a subset of those shapes c of C such that $map(c) = s$.

The set of operations are:

1. Compute $map(c)$ where c is a member of C.

2. $paint(\cdot)$ generates a visualization of a shape bucket s or a comment c.

3. $balance(\cdot, D)$ places the shape buckets S or a set of comments C' in the 2D display with dimension D. It returns the (x, y) coordinates of these objects.

4. $retrieve(\cdot)$ retrieves the content of a shape bucket s or a circle/comment c.

5. $lookup(s.c, p)$ returns `true` if c in s satisfies a user-defined *predicate* p.

6. $update(\cdot, param\_list)$ updates $s \in S$ or $s.c$ as specified in the $param\_list$ and returns the updated s or c.

In this work, we consider circle as the shape for representing a comment. We refer to the abstraction as *Several Concentric Circles* (SCC). As we shall see next, the implementation of the *map-paint-balance* abstraction is inspired by the paintings of Kandinsky. Similar to the aforementioned artworks, circles are the *motif*[4] of the SCC abstraction (**K1** in Table 1). We refer to the 2D display space as *Kandinsky Canvas* (canvas for brevity).

[4]Motif is an element or combination of elements repeated often enough in a composition to become a dominant feature in a painting.

Algorithm 1 The map algorithm

```
Require: circles - array of circles representing the comments of the post to visualize
Ensure: the root circle of each shape bucket
function MAP(circles)
    roots ← []
    circleMap ← new Map⟨commentId, circle⟩
    for each circle in circles do
        circleMap.SET(circle.comment.id, circle)
        if circle has parent then
            parentCircle ← circleMap.GET(circle.comment.parentId)
            parentCircle.children.PUSH(circle)
        else
            roots.PUSH(circle)
        end if
    end for
    return roots
end function
```

5 The SCC Operations

In this section, we describe the implementation of the operations in the SCC abstraction. Intuitively, any implementation of the operations aim to facilitate automatic *transformation* of each conversation network to an aesthetically pleasing, color-coded shape bucket and then judicious *placement* of these buckets on the *Kandinsky Canvas* to galvanize search and exploration.

5.1 Map Operation

The map operation $map(\cdot)$ maps each comment $c \in C$ to a shape bucket $s \in S$. Specifically, all comments in a conversation network are mapped to the same shape bucket. For example, all the comments in Figure 3(a) are mapped to the same shape bucket, which is visualized in Figure 3(b) using the paint operation. Hence, if there are k conversation networks associated with an anchor post, then k shape buckets are created (*i.e.,* $|S| = k$) and the comments are distributed to these buckets.

Algorithm 1 outlines the map operation. It accepts the array of circles to display on the canvas as input and returns the root circle of every shape bucket. The return value is chosen so that the buckets can then be passed to the paint operation, which is often invoked after the map operation. The time complexity of the operation is $O(|C|)$ where $|C|$ is the number of circles.

5.2 Paint Operation

Intuitively, the paint operation $paint(\cdot)$ is responsible for "painting" on the *Kandinsky Canvas*. Since the SCC data model involves two types of data, circles (*i.e.,* comments) and shape buckets, which leads to two primary goals: generating visual representations for a shape bucket and the comments it contains. We discuss them in turn.

Visual representation of a shape bucket. "Painting" a shape bucket on the canvas is inspired by Kandinsky's artworks. Each comment in a shape bucket s is represented by a colored circle where distinct colors (or color hues) are used as identifiers of the commentators within a conversation network N_s represented by the bucket (**K1** in Table 1). Color-coding these circles is a natural way to gain attention [16, 44] and visually convey different contributors and comments involved in a conversation. Replies to a comment are visualized as *concentric circles* (**K2** in Table 1). The radius of a circle is a function of the number of likes and its position in the sequence of concentricity represents its

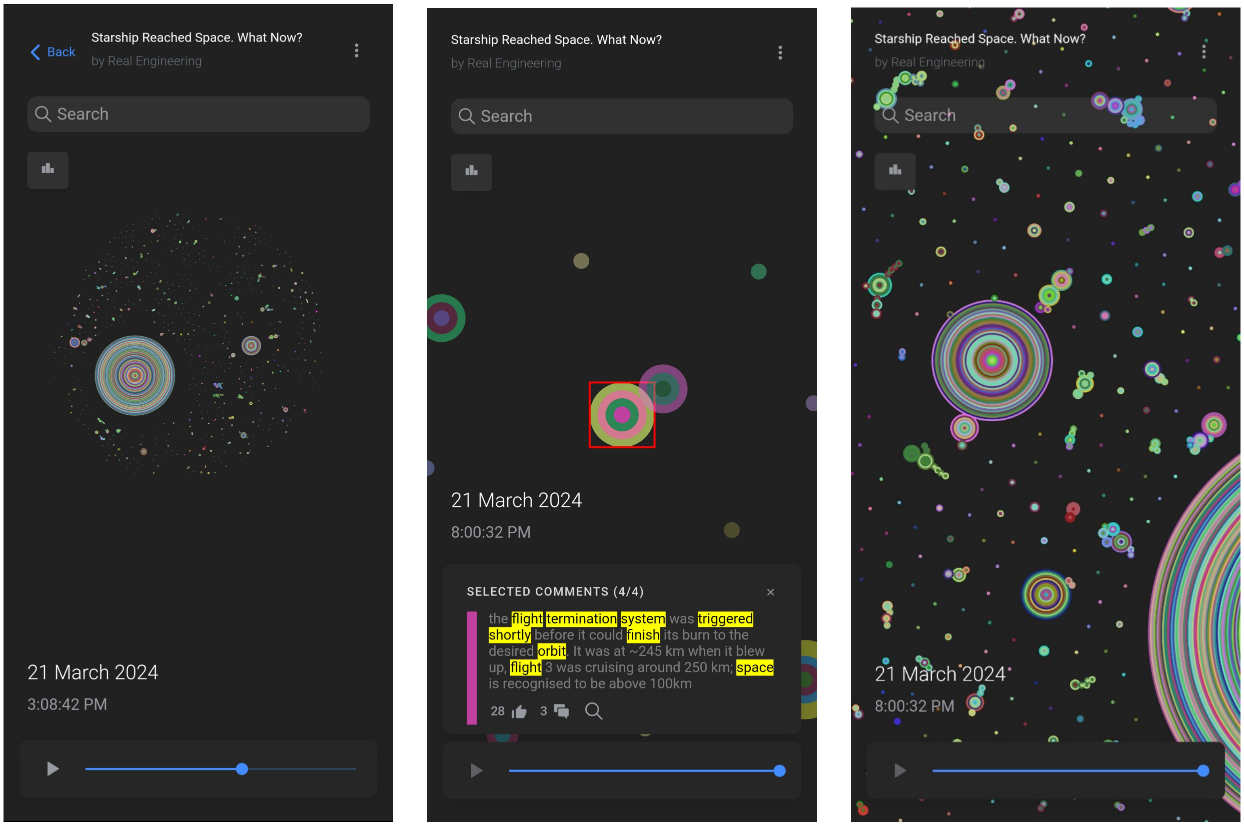

Figure 4: **[Best viewed in color]** Screenshots of Kandinsky Mobile: (left) Bird's eye view of social conversation associated with an anchor post; (middle) Comments associated with a discussion thread; (right) Examples of concentric circles and peripherals.

chronological order (ties are broken arbitrarily) in N_s. The innermost circle of a concentric circle is called a *pivot*. Intuitively, there are two types of concentric circles, *nucleus* and *peripheral*. Concentric circle resulted from an original comment and its replies is referred to as *nucleus*. In other words, in a conversation network, the root and its children form the nucleus with the root being the pivot. On the other hand, replies at depth equal to or greater than two in a conversation network (*e.g.,* replies to a reply) form separate concentric circles, referred to as *peripherals*. Peripherals are positioned on the periphery of the nucleus (**K3**).

The paint operation takes the comments associated with a shape bucket s as input and constructs the conversation network N_s. Then it generates the concentric circles by traversing N_s to visually represent the shape bucket. First, the pivot and nucleus are created and then the peripherals, if any. The size of a shape bucket is set to at least 9.2 mm by default based on the recommendation from Parhi *et al.* [45] that target sizes should be at least this value for discrete single-target tasks in *fat thumb-friendly* devices without degrading performance and preference. Note that most mobile users rely on single-touch and most use their thumbs as the only input method [5, 10]. The color codes of commentators are assigned based on the number of distinct commentators in N_s as well as the selected colors of the adjacent commentators in order to maintain visual contrast[5]. To this end, nodes in N_s are annotated with color codes that are chosen based on the commentators and colors of their parents and siblings (**K4**). These codes are then exploited during the generation of the concentric circles.

[5]Contrast in arts is a large difference between two things; for example, warm and cool colors, light and shadow. Contrasting colors and textures add excitement, emphasis and interest to a work of art.

Algorithm 2 The paint algorithm.

```
Require: buckets - the array of root circles of each shape bucket
Ensure: the nucleus concentric circle of each shape bucket
function PAINT(buckets)
    shapeBuckets ← []
    Generate radius scale from like count of all comments
    for each bucket in buckets do
        PAINTCIRCLE(bucket)
    end for
    for each bucket in buckets do
        ConcentricCircle ← PAINTCONCENTRICCIRCLE(bucket)
        shapeBuckets.PUSH(concentricCircle)
    end for
    return shapeBuckets
end function
```

For example, consider Figure 3(b), which "paints" the shape bucket associated with the conversation network in Figure 3(a). The color codes of the circles follow that of the comments. The paint operation creates the nucleus using Comments 1 (pivot), 2, 3, and 8. Note that 1 is the original post whereas 2, 3, and 8 are replies to it. Since the depths of Comments 4, 5, 6, and 7 in the network are greater or equal to 2, two peripherals are generated by the paint operation with Comments 2 and 3 being the pivots, respectively. Note that paint may recursively create peripheral-of-peripheral if the hierarchy of comments is deep (*e.g.,* Figure 4(right)). Observe that conversations with more participants result in painting of a larger and more colorful circles (**K5** in Table 1). On the other hand, the shape bucket of an original comment that has not garnered any comments degenerates to a single-color circle. Figure 4 (left) depicts results of the paint operation on the shape buckets associated with an anchor post.

Algorithm 2 outlines the paint operation. It invokes two procedures, *PaintCircle* and *PaintConcentricCircle*, to paint individual circles and concentric circles, respectively. The *PaintCircle* is a recursive function such that given a pivot circle, it paints the circle by modifying its visual properties and is recursively called on the next circle in the sequence of the bucket. The *PaintConcentricCircle* method, on the other hand, is another recursive function that paints concentric circles as a whole. Given a pivot circle, it then creates the concentric circle object by grouping this pivot with its replies. It is then recursively called to generate the peripheral concentric circles of replies that are of depth 2 with respect to this pivot circle (*i.e.,* reply to a reply). The time complexity of the operation is $O(|C|)$ where $|C|$ is the number of circles.

Visual representation of a comment. Given a comment associated with a shape bucket, topics in a textual content are highlighted in *yellow* (a primary color in Kandinsky's artworks) by paint. Colored bars are drawn on each comment item to represent the contributor (same color code) and the comment's *normalized* like count (using thickness of the bar). Figure 4 (middle) depicts an example of a comment.

Connection to visualization theories. Observe that the visual representation of shape buckets is grounded on the *expressiveness* and *effectiveness criteria* of visualization [39]. The former determines whether a graphical language can express the desired information whereas the latter identify which of these graphical language is the most effective for a given situation at exploiting the capabilities of the output medium and the human vision system. Specifically, Mackinlay [39] identified and ranked perceptual tasks for encoding quantitative, ordinal and nominal data. The paint operation is cognizant of these visualization theories and principles. Specifically, it only encodes comments that are associated with a conversation network – nothing more, nothing less. It utilizes perceptual tasks such as position, shape, color hue and area for effective visualization. Since the set of comments is nominal data, position and color hue are the most effective perceptual tasks for representing them. Furthermore, it uses the area of a (concentric) circle to capture the number of comments garnered by a conversation network. Since this is a quantitative data, area is considered as moderately effective in the ranking of perceptual tasks, which is exploited by Kandinsky Mobile. Note that higher-ranked tasks for quantitative data (*e.g.,* length, angle, slope) are not conducive for visually representing shape buckets.

5.3 Balance Operation

Given the visual representations of shape buckets and comments, the goal of the balance[6] operation *balance*$(\cdot)$ is to arrange them on the *Kandinsky Canvas* with dimension D *effectively.* It operates at two levels, the set of shape buckets and comments in a shape bucket. We discuss them in turn.

Arrangement of shape buckets. Given the set of shape buckets representing k conversation networks, we need to *effectively* place them on a mobile device interface with dimension D. This is a challenging problem since D is small and there can be many conversation networks that are large. First, the placement of the shape buckets has to be *fat thumb-friendly* [8, 10]. That is, the size of the objects and distance between them need to be palatable to fat thumb interaction. Second, it is paramount to maintain the visual clarity of the bird's-eye view of the entire discussion landscape on the canvas. These can be conflicting goals when k grows. Due to fixed D, larger k restricts the distance between the shape buckets and their sizes, adversely impacting effective data visualization. Hence, it is essential to judicially arrange the shape buckets on the canvas such that "significant" conversations are visualized effectively to engage end users and then use pinch-and-spread gesture to explore various parts of the canvas.

Given a set of visual representation of the shape buckets, the balance operation extends force-directed simulation layout and exploits fat thumb-friendly heuristics to implement it. The minimum space between elements in a mobile device is recommended to 2mm [8]. Although this is easy to achieve for icons (*e.g.,* search), there may be too many shape buckets on the canvas to impose the recommended size. Furthermore, when the rest fingers are grasping the device, the range of motion available to the thumb is more limited, thus limiting the area of the screen a user's thumb can reach. Hence, the balance operation utilizes the following three strategies:

1. Shape buckets that correspond to conversations that have garnered significant attention are placed at around center of the screen for fat thumb-friendly access [24, 25] (**K5**). In particular, touch targets in the center can be smaller - as small as 7mm, while corner target sizes must be about 12mm [25]. For example, conversations that have garnered significant attention[7] in Figure 4 (left) appear in the center of the screen.

2. A padded collision force of 30px (7.94mm) is added to the force-directed simulation layout to make the concentric circles push each other, avoiding overlaps and providing adequate space for fat-thumb support (**K6**).

3. Zooming is enabled so that one can zoom into a specific area of the display to view shape buckets in that area.

Based on the above strategies the balance operation assigns a position (*i.e.,* (x, y) co-ordinates of the pivot) to each shape bucket and returns the set of shape buckets for display on the canvas. Algorithm 3 outlines the procedure for arranging the shape buckets. The worst-case time complexity is $O(klogk + k^2)$

[6] Balance in arts is a principle of design. The arrangement of elements in a work of art (including size and number of objects) that achieves a sense of equality.

[7] In arts, this is called *focal point* – the most important part or area in a work of art.

Algorithm 3 The balance algorithm.

Require: Shape buckets S
Ensure: *none*
```
function BALANCE(concentricCircle S)
    S.EnableZoom() /* Strategy #3 */
    S.InitializeGravitionalForce() /* Strategy #1 */
    S.InitialiseRepellentForce() /* Strategy #2 */
    S.ForceSimulation()
end function
```

where k is the total number of concentric circles. It may invoke the update operation (discussed later) to modify the sizes of the concentric circles in a shape bucket, if necessary, to realize the placement of the buckets in *incremental mode*. Observe that the visual representation and arrangement of the shape buckets are aligned with the similarity, continuation, closure, proximity, and symmetry of Gestalt principles for describing visual perception [31, 47].

Arrangement of comments. Due to limited space on mobile devices, balance restricts the arrangement of a comment to fixed height adjusted as a portion of the device's length. To enhance user navigation, it is positioned at the bottom of the device to facilitate fat-thumb interaction and vertical scrolling (*e.g.,* Figure 4 (middle)).

5.4 Retrieve Operation

The retrieve operation $retrieve(\cdot)$ is designed to efficiently retrieve comments associated with a shape bucket. Given a selected shape bucket s (*e.g.,* through touch), it retrieves the content of the conversation associated with s (*e.g.,* Figure 4 (middle)) by leveraging the indexes. The current implementation supports two types of selection: (a) If one selects anywhere inside a shape bucket then the *entire* conversation network is retrieved. That is, we do not allow fine-grained selection of an individual concentric circle since it can be physically challenging in a mobile device. (b) If a peripheral is selected, then the conversation associated with it is retrieved only.

5.5 Lookup Operation

The lookup operation $lookup(\cdot)$ on a circle c (comment) in a shape bucket returns a Boolean value depending on whether the c satisfies a user-specified predicate p. Note that p can be a set of keywords (*e.g.,* "Florida") or a comment. In the case of the former, it returns true if c contains the keywords. In the latter case, it returns true if c contains comment that is *similar* to the specified comment.

5.6 Update Operation

The update operation $update(\cdot)$ updates a shape bucket on the *Kandinsky Canvas* by modifying various properties on demand. Given a shape bucket s, it may add new circles (*i.e.,* new comments in the conversation network) to s or modify properties such as color, size, and positions of existing circles in s or *highlight* it. Typically, it is invoked by the aforementioned operations. For instance, in paint the number of peripherals influences the size of the circles since we do not want the peripherals to obscure the nucleus or one another. Hence, it is necessary to judiciously set the size and positioning of them to optimize perceptibility. As more circles are added to a shape bucket, the update operation maintains these properties accordingly. Similarly, it can resize a shape bucket in response to the balance operation to optimize arrangement of shape buckets on the canvas. The lookup operation also invokes

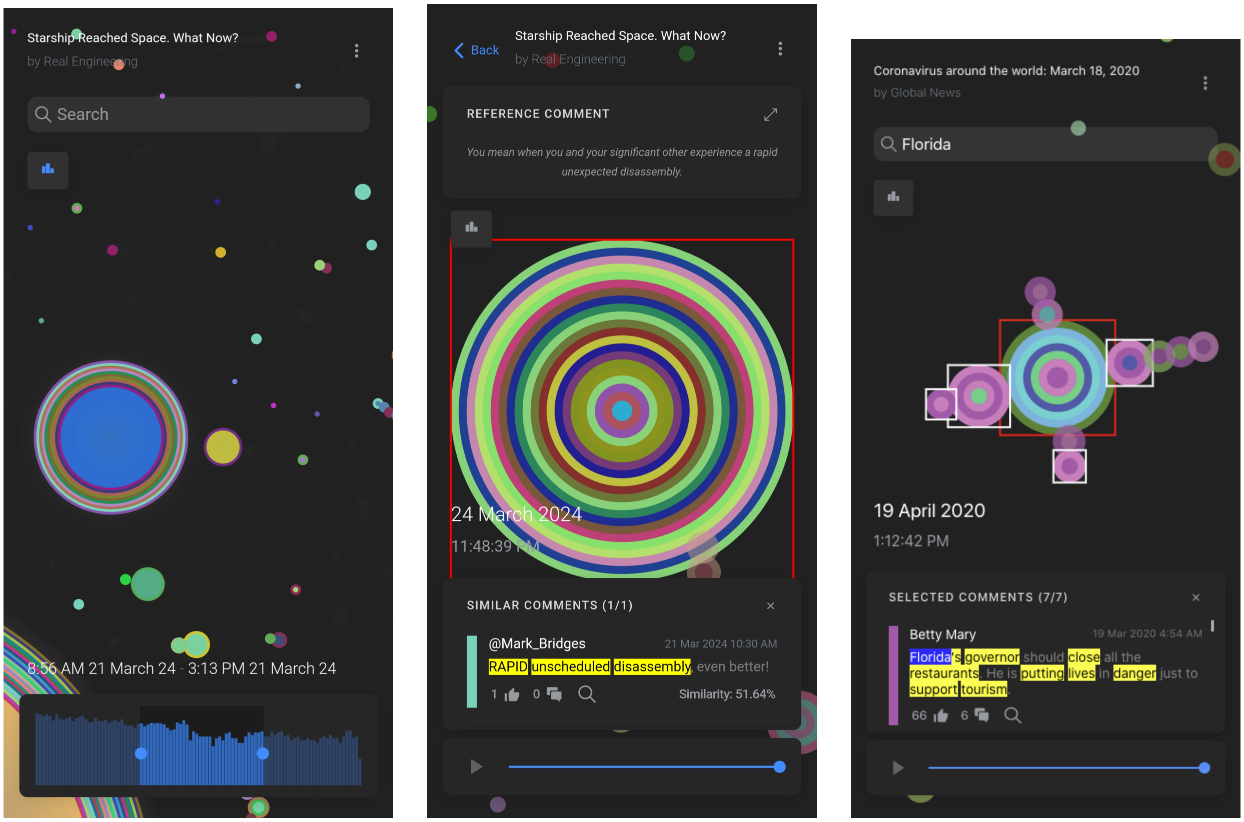

Figure 5: **[Best viewed in color]** Screenshots of Kandinsky Mobile: (left) Time travel; (middle) Similar comment search; (right) Keyword search.

it to highlight a shape bucket that satisfies the predicate by drawing a white boundary box around a circle/bucket (**K7**). An example is given in Figure 5 (right).

6 Implementing The Visualization Goals

In this section, we first briefly describe how the proposed SCC abstraction enables us to achieve the visualization goals mentioned in Section 1 with ease. We invite the reader to view the video of Kandinsky Mobile at `https://youtu.be/ay_0LpRRBQE` for an example. Next, we highlight the superiority of SCC abstraction to the popular list-based abstraction widely used in social media w.r.t. our goals.

6.1 SCC-based Visualization

Goal 1: Bird's eye view of discussion landscape. It is intuitive to visualize the discussions landscape associated with an anchor post using the SCC abstraction. Specifically, we provide two ways to generate a bird's eye view, *spectral* and *incremental*. In the *spectral mode*, the entire set of conversation networks is displayed on the canvas in one shot. This can be achieved by the following sequence of operations: [map, paint, balance]. In the *incremental mode*, the canvas is incrementally updated according to the chronological order of the comments. That is, the shape buckets are visualized incrementally over time. Note that this can be achieved by invoking the update operation with the aforementioned operations. Figure 4 (left) depicts a bird's eye view of an anchor post. The canvas has a dark background similar to *Several Circles* and *Heavy Circles* (**K8**). Observe that a user gets a fair

Table 2: Anchor posts for experimental study.

PostID	Category	Num. of comments	PostID	Category	Num. of comments
FJ1wak5zOhQ	News & Politics	63	G0WtbB-G6fc	Self-help	2564
tTBsGArSCuE	Self-help	127	_bcfxty39Cw	Science & Technology	2657
PfPelzWeSvY	Self-help	240	mY8RqxCt_Fk	Entertainment	2759
gn4S7lQ111M	Education	278	HHPX_QdI804	Education	2817
pvo6hb8rs40	Entertainment	417	bWr-DA5Wjfw	News & Politics	2891
5LpaVYxTedE	Education	465	qOOnBTaHG_Q	Education	3003
-Nm9l5qRWQU	News & Politics	602	HffWFd_6bJ0	Self-help	3069
m6qFCdHvYuI	Science & Technology	747	TlSOMfDX-yY	Science & Technology	3228
NFLjeyd2M0k	Science & Technology	819	2SkLBP0QtSQ	Entertainment	3307
g1pb2aK2we4	Education	961	eaFyTIsowqY	Science & Technology	3444
RThSD70hiTg	Self-help	1059	Ra4W_aztfHA	News & Politics	3545
oRt5NPY7yYA	Science & Technology	1163	P6FORpg0KVo	Self-help	3612
DAKRcS_XeEI	News & Politics	1211	B9SptdjpJBQ	Self-help	3696
EFRUL7vKdU8	Science & Technology	1332	qu6GMpQfyUU	Education	3799
pA5hNeJe1_I	Entertainment	1418	sm18t0m0hRc	Entertainment	3938
VsBSCj2oqwU	Entertainment	1458	MA7TQNExbRg	Education	4079
HFYv-rk4v9Y	News & Politics	1615	tPuOeg4zGQQ	News & Politics	4217
Y2M2Dgmm958	Self-help	1701	xX96xng7sAE	Education	4306
BaptBhVD1r0	News & Politics	1775	YMZcp0EQO2s	Self-help	4356
wl_s54MBsmY	Entertainment	1910	3kwDVw0u4Kw	Science & Technology	4535
c-jALYCMOz0	Entertainment	1998	-CcQ4jKn8aE	News & Politics	4616
rD5goS69LT4	Self-help	2095	sXJg9J81acY	Science & Technology	4685
8AaPHNlrwFU	Entertainment	2159	0r2x7G0hwCw	Education	4768
M2pzKxMNKOM	Education	2290	fkbZ60JXXB8	Entertainment	4921
1CHt6Yo6sVE	News & Politics	2430	kRzgCylePjk	Science & Technology	4984

Table 3: Distribution of comments in anchor posts used in experiments.

Number of Comments	Number of Samples
1-1000	10
1001-2000	10
2001-3000	10
3001-4000	10
4001-5000	10

opportunity to *quickly* select *any* of the shape buckets to peruse the corresponding conversation instead of getting steered by ML-based algorithms.

Goal 2: On-demand visualization of conversation. Given the constructed *Kandinsky Canvas*, one can explore the canvas space and select any shape bucket or peripheral to view the corresponding comments. This is achieved by the following sequence of operations on the selected bucket or peripheral: [retrieve, paint, balance].

Goal 3: Canvas search. The *canvas search problem* takes the user-specified keywords (resp. comment) as input and returns an updated canvas highlighting shape buckets containing the keywords (resp. similar comments). The sequence of operations to realize this is [lookup, update]. For each circle in a shape bucket, the lookup operation returns true if it contains the user-specified keyword or similar comment. Then the update operation surrounds the corresponding shape bucket or peripheral with a white bounding box (**K7** in Table 1). Figure 5 (right) depicts an example of keyword search. For similar comment search, it also lowers the color intensity of the shape buckets (*i.e.,* modifies the color properties of the circles) that do not contain similar comments to make them less visible (**K9**). Selecting one of the highlighted shape buckets retrieves and displays its comments using [retrieve, paint, update, balance] where the update operation highlights the matching keywords in yellow. Figure 5 (middle) depicts an example of similar comment search. The reference comment is shown at the top and similar comment is displayed below along with the shape bucket.

Goal 4: Time travel. Lastly, the SCC operations enable time travel to visualize the discussion landscape (*i.e., Kandinsky Canvas*) at a particular time period (referred to as *Spectrum Filter* in

Table 4: Experimental environment.

Machine	Xiaomi Poco X3 Pro
Operating System / Version	Android OS / Version 12
Processor	Qualcomm Snapdragon 860
Memory	8GB
Internet Connection / Speed	Wi-Fi Connection / 76.8Mbps

Kandinsky Mobile). Given a user-specified time duration $[t_1, t_2]$ and the canvas, comments that are not published within the specified time frame are blurred out (**K9**). This can be achieved by invoking update to "repaint". That is, change the color property of the circles in each shape bucket that are not published between t_1 and t_2. Figure 5 (left) depicts an example.

6.2 List-based vs SCC Abstractions

All major commercial social media sites visualize social discussions in form of scrollable and expandable list. In particular, comments that are automatically exposed to a user are selected by secret-sauce machine learning-based algorithms that aim to optimize engagement and profit. Such abstraction is suboptimal in providing a bird's eye view of the discussion landscape. In contrast, the SCC abstraction enables end users to visualize and explore the entire landscape with ease. Importantly, it gives control back to them to decide which conversation network they would like to engage with instead of feeding them conversations selected by opaque ML algorithms. Furthermore, SCC has a richer set of operations that support aesthetically pleasing visualization and on-demand search and exploration of the social discussions. Lastly, SCC can be easily used to create effective, aesthetically pleasing visualization for solutions to important problems. For example, one can use any state-of-the-art technique to detect sentiments [34] of the discussions and then invoke the update operation to modify the properties of shape buckets to visualize different sentiments in an conversation. Similarly, it can emphasize the comments of opinion leaders in an anchor post by dimming the colors of circles associated with non-leaders. Such visualization is not easily realizable in existing list-based abstraction.

7 Performance Study

Kandinsky Mobile is built using *Ionic 4*, a framework for building hybrid mobile applications. It currently supports anchor posts in *YouTube*. In this section, we report preliminary performance results of Kandinsky Mobile. We also report an initial small-scale user study to report the potential usefulness of the abstract art-inspired SCC abstraction. The codebase is available at `https://github.com/chun-leong/kandinsky-mobile`.

7.1 Efficiency and Scalability

Experimental Setup. We selected 50 *YouTube* anchor posts under five different categories (news & politics, self-help, education, entertainment, and science & technology) containing up to around 5000 comments. Each category contains 10 anchor posts. Table 2 reports the *postid*, *category*, and the *number of comments* of these anchor posts. Table 3 shows the distribution of the comments.

Table 4 reports the system environment for our experimental study. We ensure that no other application processes are running on the machine. Also, we ensure that any Kandinsky Mobile application data is deleted prior to running each experiment.

An experiment on each anchor post is executed 3 times to obtain the average runtime performance. In order to measure end-to-end performance of Kandinsky Mobile, instead of measuring the performance

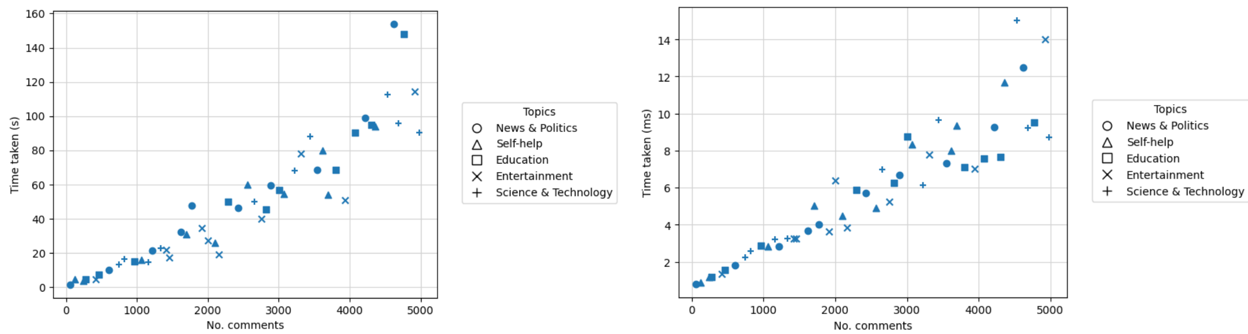

Figure 6: Performance of the extraction stage: (left) Response time; (right) Storage time for post metadata.

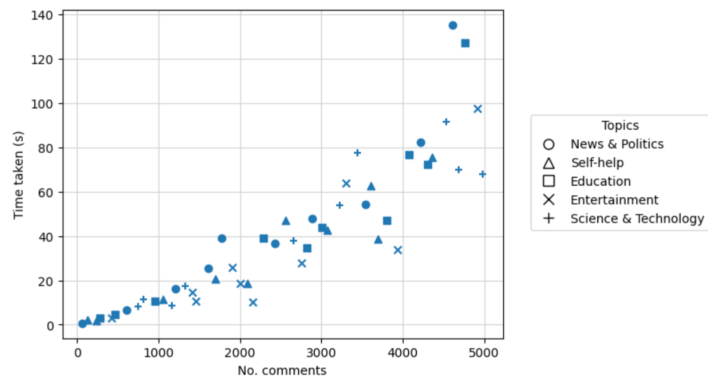

Figure 7: Performance of the representation stage.

of individual SCC operations, we measure the performances of the three stages associated with any anchor post: *extraction*, *representation*, and *canvas construction*. We elaborate on them in the turn.

Extraction. The *extraction* stage in Kandinsky Mobile is responsible for retrieving the post and comment data as well as converting them into a form usable by the other services in the Kandinsky Mobile application. The time taken for the extraction stage includes fetching the data through the social media platform's data API, parsing the information, and saving the data to the persistent storage. Note that these are the processes needed to be completed before the anchor post data can be utilized by Kandinsky Mobile.

Figure 6 (left) plots the response time of the extraction stage. We can make the following observations. First, it grows linearly with the number of comments. Second, given that a large majority of anchor posts typically draw lesser than 3000 comments, the extraction stage can be completed within a minute for these posts.

Note that majority of the extraction time is dominated by the data retrieval step through the *YouTube* API. Figure 6 (right) reports the time required to save each post's metadata to the persistent storage. Observe that it grows linearly with the number of comments and can be completed efficiently.

Representation. The *representation* stage in Kandinsky Mobile performs the analysis of comment data. Specifically, in this stage topic modeling is performed, which has a time complexity of $O(m)$ where where m is the total number of words in all comments. The results obtained from performing the representation stage are used in the exploration of the conversation landscape. Figure 7 plots the results. Observe that for majority of the anchor posts, the representation stage can be performed within a minute.

Canvas Construction. This stage is responsible for representing and visualizing comments on the

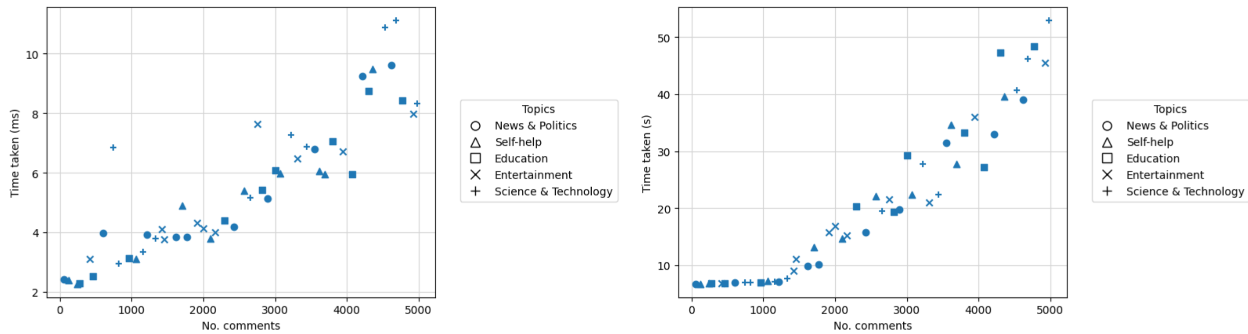

Figure 8: Performance of the canvas construction stage: (left) Canvas preparation; (right) Canvas drawing.

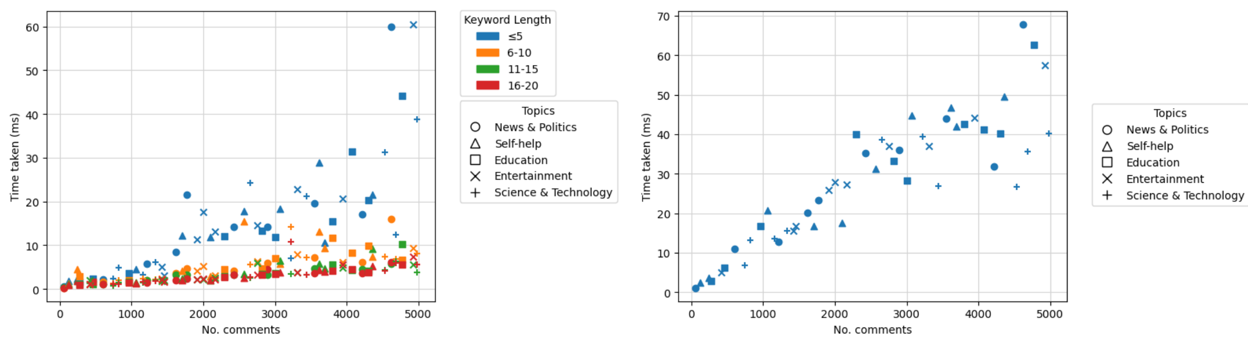

Figure 9: **[Best viewed in color]** Performance of the visualization goals: (left) Keyword search; (right) Similar comment search.

canvas. In other words, it performs the *Canvas Construction Service*, which invokes the map, paint, and balance operators. This stage is the final stage of visualizing the conversation landscape of a post. After which, users may carry out their exploration of the landscape as they desire.

The canvas construction stage can be decomposed into two sub-stages, *canvas preparation* and *canvas drawing*. *Canvas preparation* includes generating the circle entities from the comments. Hence, the map operation is performed in this sub-stage. *Canvas drawing* includes performing the paint and balance operations on the canvas entities. Figure 8 (left) reports the performance of the canvas preparation stage. Observe that the running time increases linearly to the number of comments. It is very efficient and can finish processing the largest number of comments within 12 msec.

The performance of canvas drawing is reported in Figure 8 (right). Observe that we can generate the canvas within a minute for posts with around 5000 comments. For majority of posts with less than 3000 comments, it can generate the canvas within 20 seconds.

Performance of our visualization goals. Lastly, we report the performance of some of the visualization goals described in Section 6. The bird's eye view generation time (Goal 1) is similar to the canvas construction time reported above. The running time for canvas search (Goal 3) and time travel (Goal 4) is less than 70 msec and hence highly interactive. Figure 9 plots the performance of keyword search for different keyword set size and similar comment search. Observe that they grow linearly with the number of comments.

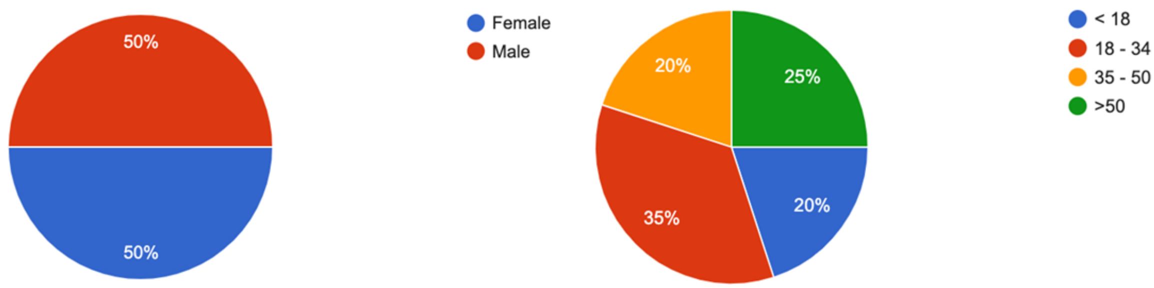

Figure 10: **[Best viewed in color]** Profile of participants in the user study.

7.2 User Study

Next, we undertake a small-scale user study to investigate the usefulness and usability of Kandinsky Mobile w.r.t. the four visualization goals. Specifically, we investigate the following research questions:

RQ1: How is the experience, engagement, and usability of Kandinsky Mobile as an end-to-end mobile app?

RQ2: How effective is Kandinsky Mobile w.r.t. the four visualization goals?

RQ3: How does the SCC-based abstraction in Kandinsky Mobile compare with the popular list-based abstraction?

Note that we do not focus on how closely the design of Kandinsky Mobile is aligned to the artworks of Kandinsky, as it would be unrealistic to expect our participants to possess expertise in his body of work. We first present the setup for our user study and then highlight the key results.

Participants profile. 20 unpaid volunteers initially participated in the study in accordance to HCI research that recommends at least 10 participants [19, 36]. The gender and age distributions of the participants are plotted in Figure 10. As Kandinsky Mobile is intended for users with normal vision and familiarity with social interactions on social media platforms, we collected feedback from them regarding their experiences with social media posts, usage of mobile devices, and experience in using digital technology. All subjects reported to have normal vision. Figure 11 depicts the survey form and the results on a Likert scale of 1-10 with higher the score the greater is the agreement with the statement. Observe that majority of the users use their mobile phones a lot, regularly interact with social media posts, and are savvy with digital technology. Specifically, three participants (two female, one male) indicated limited exposure to social media posts. Consequently, we replaced them with three individuals of the same gender and demographic background who had adequate experience with social conversations, to better isolate the impact of Kandinsky Mobile from confounding factors such as limited exposure to social media discourse.

Tasks. Next, we presented a scripted tutorial on how to use Kandinsky Mobile. All participants were then requested to use their own mobile phones to use it and explore social discussions in *YouTube*. The specific anchor posts were chosen by them and they were instructed to perform the following tasks using *both* list-based and SCC-based visualizations for each post in random order:

(a) find conversations that have attracted significant discussions (related to Goal 1);

(b) select conversations of interest and read them (related to Goal 2);

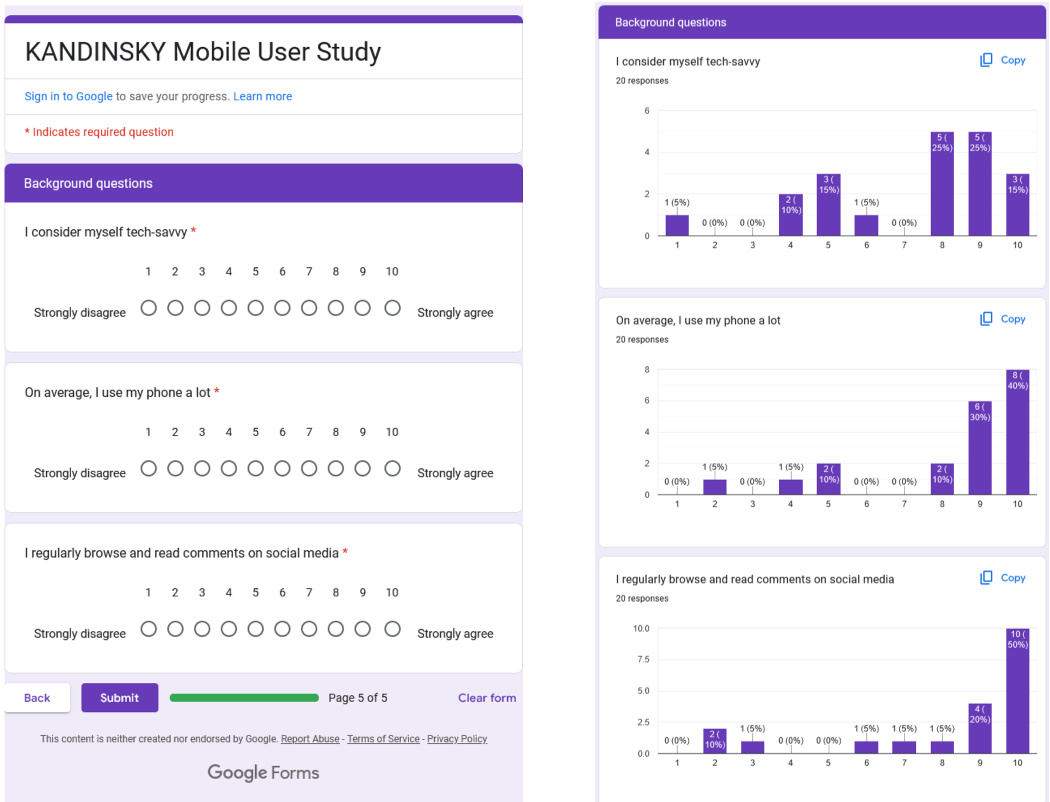

Figure 11: Experience of participants of the user study.

(c) find discussions containing comments with some keywords-of-interest (related to Goal 3);

(d) find similar comments to a chosen comment (related to Goal 3); and

(e) find all discussions that have occurred up to a particular timepoint in the history (related to Goal 4).

Finally, they were asked to answer a survey. Each subject gave a rating in the Likert scale of 1-10 for each question in the survey with higher the score the greater is the agreement with it. On average each participant performed the above tasks on 7.2 anchor posts.

Quantitative Results. We first seek feedback on the overall experience, engagement, and usability of Kandinsky Mobile (**RQ1**). Figure 12 reports the results. Most of them gave positive ratings. Next, we drill down to gather feedback on how effectively Kandinsky Mobile supports the previously stated goals (**RQ2**). The results are illustrated in Figure 13. Notably, participants generally rated its usefulness positively across all four goals, with the lowest rating being 7. Lastly, we requested the subjects to compare their experience of abstract art-inspired visualization with the prevailing list-based visualization of social conversations (**RQ3**). Figure 14 reports the results. While all subjects found the visualization superior to list-based visualization, they found Kandinsky Mobile laggy. This is due to the time taken to construct the *Kandinsky Canvas* as reported earlier.

Subjective Feedback. We also asked the participants to provide unstructured feedback on what they liked (resp. disliked) about Kandinsky Mobile. Table 5 reports some representative feedback. Observe that the subjects likes the usability and aesthetics of the interface and found it innovative. Specifically, they find the canvas "beautiful", highlighting the aesthetic appeal of Kandinsky Mobile. On the other

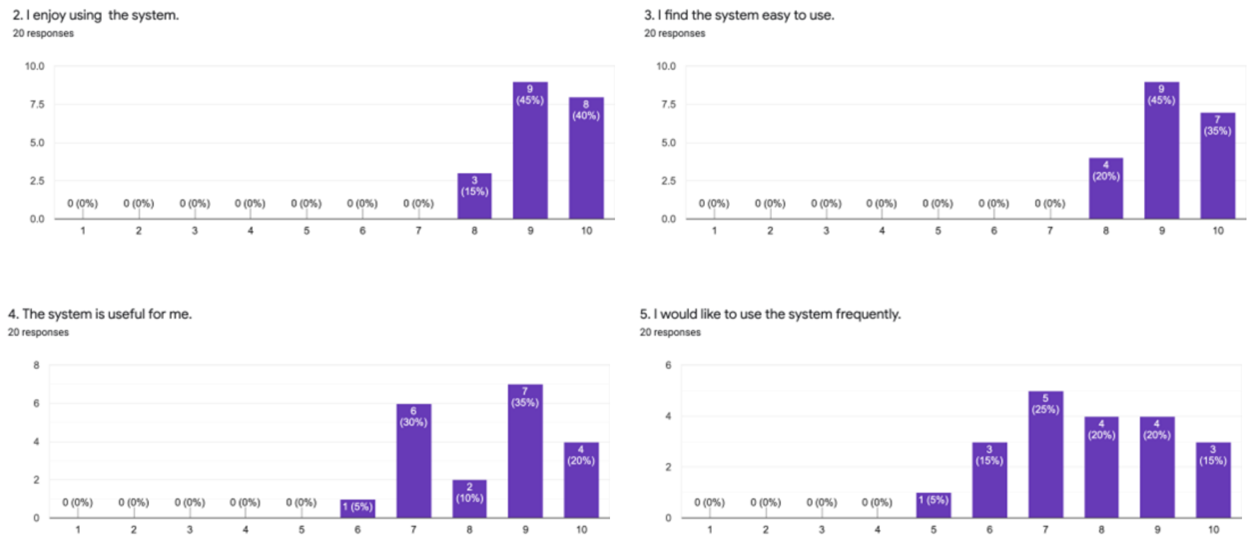

Figure 12: Overall experience with Kandinsky Mobile (**RQ1**).

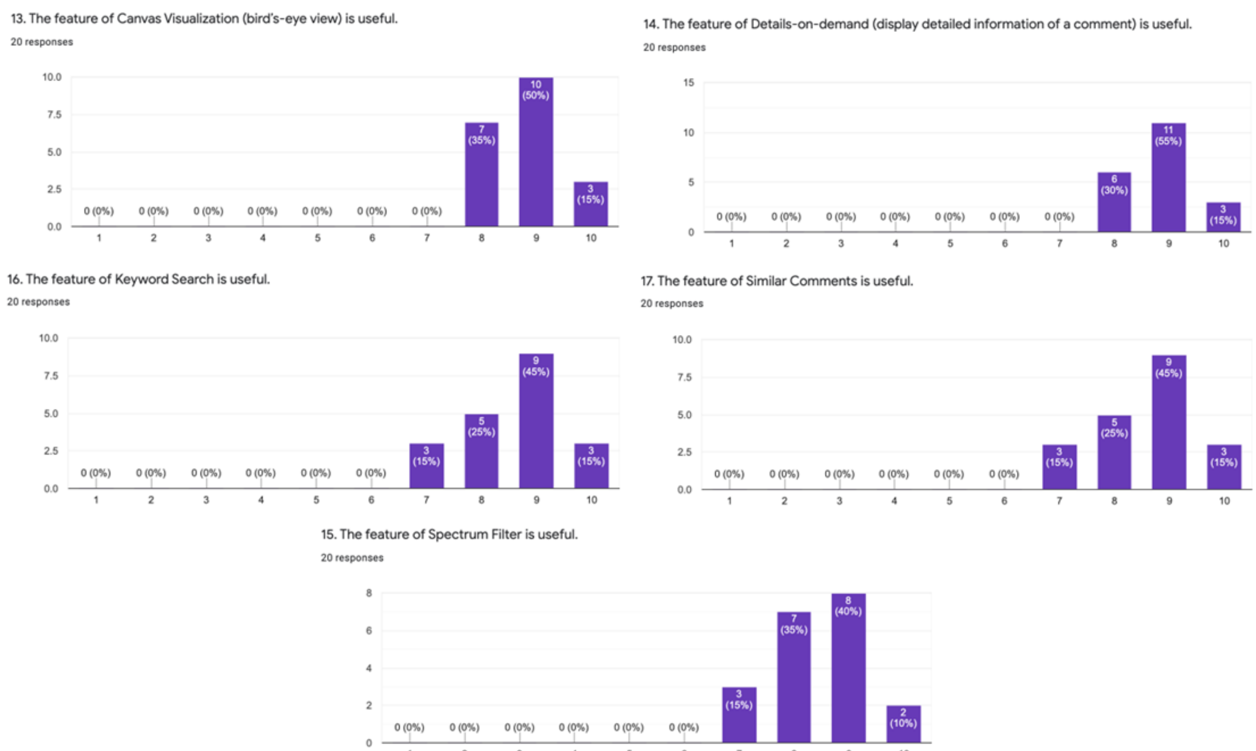

Figure 13: Usefulness of Kandinsky Mobile for supporting the four goals (**RQ2**): (Top) Bird's eye view of the conversation landscape and on-demand visualization of conversation. (Middle) Interactive search; (Bottom) Time travel.

hand, they commented on the delay in canvas construction and pointed out that the current version is limited to the *YouTube* platform only. The participants would like the current framework to be extended

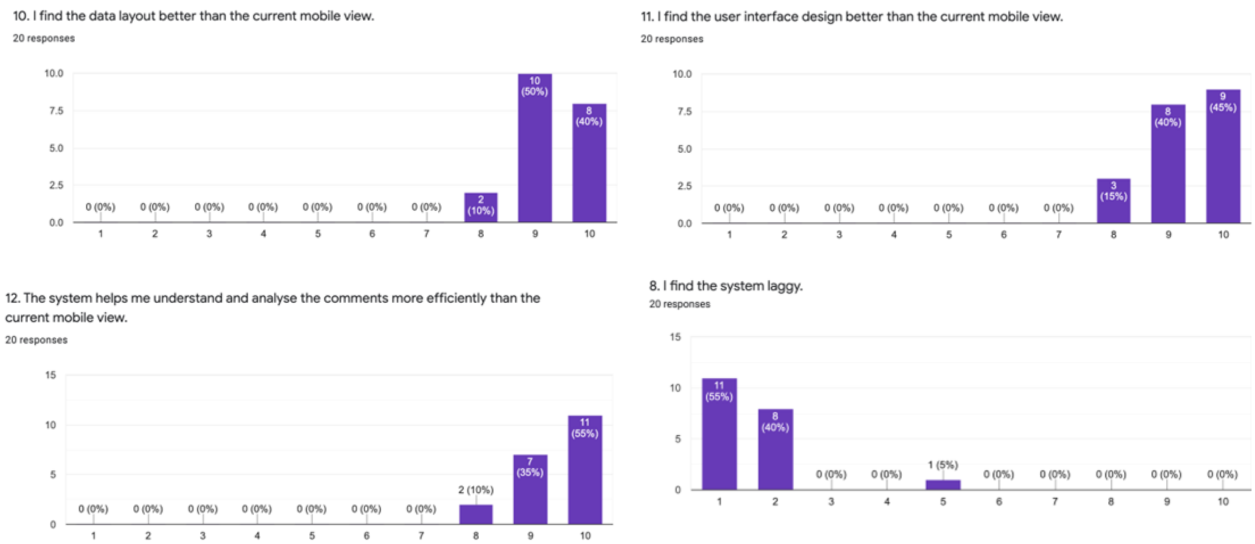

Figure 14: Comparison of abstract art-inspired visualization and list-based visualization (**RQ3**).

Table 5: Subjective feedback.

Question	Response
What do you like most about Kandinsky Mobile?	“easy to use”, “the features and functions are very useful”, “the UI looks nice”, “Nice design”, “Innovative visualization of displaying comments”, “Cool UI”, “Beautiful!”
What do you like least about Kandinsky Mobile?	“Cannot export information”, “Only supports Youtube”, “Time taken to create the canvas can be long for conversations with many comments”
What other features do you think the system should provide?	“Export data”, “Support for more platforms such as fb, X, insta”, “Sentiment analysis”, “Comment summary”, “Support other languages such as Chinese”, “Add to Google Play Store”

to other social media platforms and languages as well as support vertical applications such as sentiment analysis and comment summarization.

8 Reflections & Future Work

This paper introduces a novel visualization paradigm called *Several Concentric Circles* (SCC), which draws inspiration from the abstract art of Wassily Kandinsky and integrates it with digital technology to enable intuitive, visually engaging, and effective exploration of social discussions on mobile platforms. Unlike conventional list-based interfaces that dominate social media applications, SCC emphasizes aesthetic composition and spatial metaphors to represent conversational dynamics in a more immersive and meaningful way. At the heart of this paradigm is a working prototype we developed, named Kandinsky Mobile, which leverages the expressive potential of abstract forms and spatial layout to overcome the cognitive and navigational limitations posed by traditional linear-scrolling interfaces. To the best of our knowledge, this fusion of abstract artistic principles with mobile interaction design for the purpose of social discussion visualization represents a previously unexplored direction, first addressed in our recent demonstration of Kandinsky Mobile [22].

Enhancing Kandinsky Mobile. Our preliminary performance study demonstrates the superior potential of the SCC abstraction for visualizing social conversations. Nevertheless, there are several opportunities for improvement. First, there are opportunities to enhance the efficiency and scalability of

Kandinsky Mobile as highlighted by the participants in our user study. Although the interactive search and time travel can be undertaken efficiently, the performance of canvas construction can be further improved. In this context, while data retrieval time is the primary bottleneck during the extraction stage, both the representation and canvas drawing stages can take up to a minute for very large conversations. Therefore, there is potential for improvement by leveraging parallelism, as the k shape buckets can be potentially "painted" concurrently. Additionally, although the results of our initial user study are encouraging, we plan to conduct a more extensive study with a larger group of participants to gain deeper insights into the practical impact of Kandinsky Mobile.

Second, although the subjects in our user study generally find the SCC-based visualization aesthetically pleasing (Table 5), it is important to undertake a rigorous and systematic analysis of aesthetics of Kandinsky Mobile. In this context, we aim to investigate how color theories can inform the visual representation of a large user base. While Kandinsky's artworks typically feature only a few colored circles arranged on the canvas, Kandinsky Mobile employs a much broader color palette to accommodate the number of users.

Third, in the current framework we assume the target audience to be users with normal vision. It is interesting to explore how we can extend SCC to support individuals suffering from color blindness. Furthermore, as mentioned by some participants, it is important to integrate Kandinsky Mobile with additional social media platforms.

Lastly, anchor posts that garner thousands of comments may demand a 3D visualization of the canvas for more effectiveness. The intuition is to show "significant" conversations first and enable a user to "travel" through the depth of the canvas to explore less significant conversations. Solution to this problem demands marrying ideas from 3D mobile interactions [18, 21] with visual art-based visualization schemes.

Beyond Social Conversations. Although we focus on visualizing social media discussions, the underlying principles of SCC are broadly applicable and can be readily extended to other domains characterized by hierarchical or threaded interactions. One such example is online learning environments, such as discussion boards used in MOOCs or university learning management systems, where student-instructor and peer-to-peer exchanges often form complex, multi-level conversation trees. By adapting SCC to these contexts, it is possible to enhance user engagement, improve navigability, and provide a more aesthetically enriching and cognitively supportive interface for exploring educational dialogues. Beyond discussion boards, a recent study [37] took inspirations from SCC to visualize student feedback.

An intriguing direction for future exploration is the extension of the SCC abstraction to support the aesthetic and effective visualization of large-scale networks. Dunne and Shneiderman [15] proposed a compelling technique in this space, where common network motifs (*fans*, *connectors*, *cliques*) are identified and replaced with compact, visually distinct glyphs. This approach can be viewed as an instantiation of the SCC paradigm, where shapes are circles and lines representing nodes and edges, respectively, and are grouped into higher-level shape buckets that correspond to glyphs. Consequently, the map-paint-balance operations "paint" the glyph-based network on a display space by carefully laying out the shapes and shape buckets (*i.e.,* nodes, edges, and glyphs).

Beyond network visualization, SCC also presents opportunities as a flexible substrate for a variety of vertical applications related to social discussion analysis. For instance, sentiment analysis results could be visually encoded by leveraging color, texture, or opacity of circles–enabling users to intuitively perceive emotional dynamics within complex social conversations. Similarly, it can be the substrate for effective visualization and exploration of online hostility [41].

Acknowledgments. We would like to acknowledge Wen Shan Yap, a professional artist in Singapore for sharing his insights on abstract art and works of Wassily Kandinsky.

References

[1] Circles in a Circle (1923). https://www.philamuseum.org/collection/object/51019.

[2] Several Circles (1926). https://www.guggenheim.org/artwork/1992.

[3] Heavy Circles (1927). https://www.nortonsimon.org/art/detail/P.1953.216/.

[4] Mid Journey. https://www.midjourney.com/home.

[5] Mobile Usability Testing: Fat Finger Syndrome. https://ux247.com/mobile-usability-testing-fat-finger-syndrome/.

[6] Stable Diffusion Model. https://github.com/AUTOMATIC1111/stable-diffusion-webui.

[7] Survey of Commenters and Comment Readers. https://mediaengagement.org/research/survey-of-commenters-and-comment-readers/#:~:text=People%20Read%20Comments%20to%20Learn,or%20amused%20by%20others'%20comments.

[8] UX for Mobile: The Rise of Fat-Finger Design. https://www.outsystems.com/blog/posts/ux-for-mobile-fat-finger-design/.

[9] D. Archambault, T. Munzner, D. Auber. GrouseFlocks: Steerable Exploration of Graph Hierarchy Space. *IEEE Trans. Vis. Comput. Graph.*, 14(4): 900-913, 2008.

[10] S. Boring, D. Ledo, et al. The Fat Thumb. *In MobileHCI*, 2012.

[11] K. Chanthatrojwong, S. S. Bhowmick, B. Choi. LICS: Towards Theory-Informed Effective Visual Abstraction of Property Graph Schemas. *In ACM SIGMOD*, 2025.

[12] K. N. Diaz. Kandinsky's Nonrepresentational Art: Finding Meaning Through Aesthetic Experience. *ETD Collection for University of Texas, El Paso*. AAI1494343, https://scholarworks.utep.edu/dissertations/AAI1494343, 2011.

[13] A. Dix, J. Finlay, G. Abowd, R. Beale. Human-computer Interaction, 2nd edn, *Pearson Education Ltd*, Harlow, 1998.

[14] J. Donath, K. Karahalios, F. B. Viégas. Visualizing Conversation. *In HICSS*, 1999.

[15] C. Dunne, B. Shneiderman. Motif Simplification: Improving Network Visualization Readability with Fan, Connector, and Clique Glyphs. *In CHI*, 2013.

[16] M. A. Dzulkifli, M. F. Mustafar. The Influence of Colour on Memory Performance: A Review. *Malays J Med Sci.*, 20(2): 3-9, 2013.

[17] B. Engdahl, M. Köksal, G. Marsden. Using Treemaps to Visualize Threaded Discussion Forums on PDAs. *In CHI*, 2005.

[18] D. Fiorella, A. Sanna, F. Lamberti. Multi-touch User Interface Evaluation for 3D Object Manipulation on Mobile Devices. *J. Multimodal User Interfaces*, 4(1): 3-10, 2010.

[19] L. Faulkner. Beyond the Five-user Assumption: Benefits of Increased Sample Sizes in Usability Testing. *Behavior Research Methods, Instruments, & Computers*, 35(3), 2003.

[20] V. Giller, R. Melcher, J. Schrammel, R. Sefelin, M. Tscheligi. Usability Evaluations for Multi-device Application Development Three Example Studies. *In the Proceedings of the International Conference on Mobile Human-Computer Interaction*, Springer, Berlin, Heidelberg, 2003.

[21] E. S. Goh, M. S. Sunar, A. W. Ismail. 3D Object Manipulation Techniques in Handheld Mobile Augmented Reality Interface: A Review. *IEEE Access*, 7: 40581-40601, 2019.

[22] C. C. F. Gumtang, S. S. Bhowmick. Kandinsky Mobile: Abstract Art-Inspired Interactive Visualization of Social Discussions on Mobile Devices. *In MM*, Interactive Arts Track, 2021.

[23] S. Havre, E. Hetzler, P. Whitney, and L. Nowell. ThemeRiver: Visualizing Thematic Changes in Large Document Collections. *IEEE Transactions on Visualization and Computer Graphics*, 8, 1, 2002.

[24] S. Hoober. Design for Fingers, Touch, and People, Part 1. Retrieved from `https://www.uxmatters.com/mt/archives/2017/03/design-for-fingers-touch-and-people-part-1.php`, March 2017.

[25] S. Hoober. Design for Fingers, Touch and People, Part 3. Retrieved from `https://www.uxmatters.com/mt/archives/2017/07/design-for-fingers-touch-and-people-part-3.php`, July 2017.

[26] S. Hoober, E. Berkman. Designing Mobile Interfaces. *OReilly*, Beijing, 2012.

[27] E. Hoque, G. Carenini. ConVis: A Visual Text Analytic System for Exploring Blog Conversations. *Computer Graphics Forum*, 33, 3, 2014.

[28] I. Hossain, S. Puppala, Md. J. Alam, S. Talukder, Z. Talukder. A Visual Approach to Tracking Emotional Sentiment Dynamics in Social Network Commentaries. *In ICWSM*, 2024.

[29] K. Y. Huang. Challenges in Human-computer Interaction Design for Mobile Devices. *In Proc. of the World Congress on Engineering and Computer Science*, 1(6), 2009.

[30] W. Huang, P. Eades, S.H. Hong. Measuring Effectiveness of Graph Visualizations: A Cognitive Load Perspective. *Inf. Vis.*, 8(3), 2009.

[31] J. Johnson. Designing With The Mind In Mind: Simple Guide to Understanding User Interface Design Guidelines. *Morgan Kaufmann*, 2020.

[32] W. Kandinsky. Concerning the Spiritual in Art. *Dover Publications*, *Inc,* 1977.

[33] W. Kandinsky. Point and Line to Plane. *Courier Corporation,* 1979.

[34] W. M. Kouadri, M. Ouziri, S. Benbernou, K. Echihabi, T. Palpanas, I. B. Amor. Quality of Sentiment Analysis Tools: The Reasons of Inconsistency. *PVLDB*, 14(4), 2020.

[35] D. Lauer. Facebook's Ethical Failures are not Accidental; They are Part of the Business Model. *AI and Ethics 1*, 395-403, 2021.

[36] J. Lazar, J.H. Feng, H. Hochheiser. Research Methods in Human-Computer Interaction. John Wiley & Sons, 2010.

[37] F. Lin, C. Li, W. Ying R. Lim, Y.-H. Lee. Empower Instructors with Actionable Insights: Mine and Visualize Student Written Feedback for Instructors' Reflection. *Comput. Educ. Artif. Intell.*, 8, 2025.

[38] C. Lui, S. S. Bhowmick, A. Jatowt. KANDINSKY: Abstract Art-Inspired Visualization of Social Discussions. *In SIGIR*, 2019.

[39] J. Mackinlay. Automating the Design of Graphical Presentations of Relational Information. *ACM Trans. on Graphics*, 5(2), 1986.

[40] J. Ma, S. S. Bhowmick, B. Choi, L. Tay. Theories and Principles Matter: Towards Visually Appealing and Effective Abstraction of Property Graph Queries. *In SIGMOD*, 2023.

[41] D. Murthy, S. Sharma. Visualizing YouTube's Comment Space: Online Hostility as a Networked Phenomena. *New Media Soc.* 21(1), 2019.

[42] C. Nobre, M. D. Meyer, M. Streit, A. Lex. The State of the Art in Visualizing Multivariate Networks. *Comput. Graph. Forum*, 38(3): 807-832, 2019.

[43] M. Noirhomme-Fraiture, F. Randolet, L. Chittaro, G. Custinne. Data visualizations on Small and Very Small Screens. *In Proceedings of the 11th International Symposium on Applied Stochastic Models and Data Analysis*, 2005.

[44] Y. Pan. Attentional Capture by Working Memory Contents. *Can J Exp Psychol.* 64(2), 2012.

[45] P. Parhi, A. K. Karlson, B. B. Bederson. Target Size Study for One-handed Thumb use on Small Touchscreen Devices. *In MobileHCI*, 2006.

[46] D. Raneburger, D. Alonso-Ríos, R. Popp, H. Kaindl, J. Falb. A User Study with GUIs Tailored for Smartphones. *In IFIP Conference on Human-Computer Interaction*, Springer, Berlin, Heidelberg, 2013.

[47] A. S. Reber. Gestalt Psychology. *The Penguin Dictionary of Psychology*, Viking, 1985.

[48] C.-Y. Sung, X.-Y. Huang, Y. Shen, F.-Y. Cherng, W.-C. Lin, H.-C. Wang. ToPIN: A Visual Analysis Tool for Time-anchored Comments in Online Educational Videos. *In CHI Extended Abstracts*, 2016.

[49] T. Tanaka, M. Chosokabe, K. Tanimoto, S. Tsuchiya. Visualization of the Social Atmosphere Using Comments on News Sites. GDN 2022: 105-114

[50] E. Z. Victorelli, et al. Understanding Human-data Interaction: Literature Review and Recommendations for Design. *Int. J. Hum. Comput. Stud.*, 134: 13-32, 2020.

[51] M. Wattenberg, D. Millen. Conversation Thumbnails for Large-scale Discussions. *In CHI Extended Abstracts*, 2003.

[52] R. Xiong, J. Donath. PeopleGarden: Creating Data Portraits for Users. *In UIST*, 1999.

[53] V. Yoghourdjian, et al. Exploring the Limits of Complexity: A Survey of Empirical Studies on Graph Visualization. *Visual Informatics* 2, 2018.

[54] V. Yoghourdjian, T. Dwyer, K. Klein, K. Marriott, M. Wybrow. Graph Thumbnails: Identifying and Comparing Multiple Graphs at a Glance. *IEEE Trans. Vis. Comput. Graph.*, 24(12): 3081-3095, 2018.

[55] A. L. Zhou, Y.-A. Wang, W. Wu, K. Zhang. Human Aesthetic Preference-Based Large Text-to-Image Model Personalization: Kandinsky Generation as an Example. *In CVPR*, 2024.