Photosensitive Accessibility for Interactive Data Visualizations

Laura South and Michelle A. Borkin

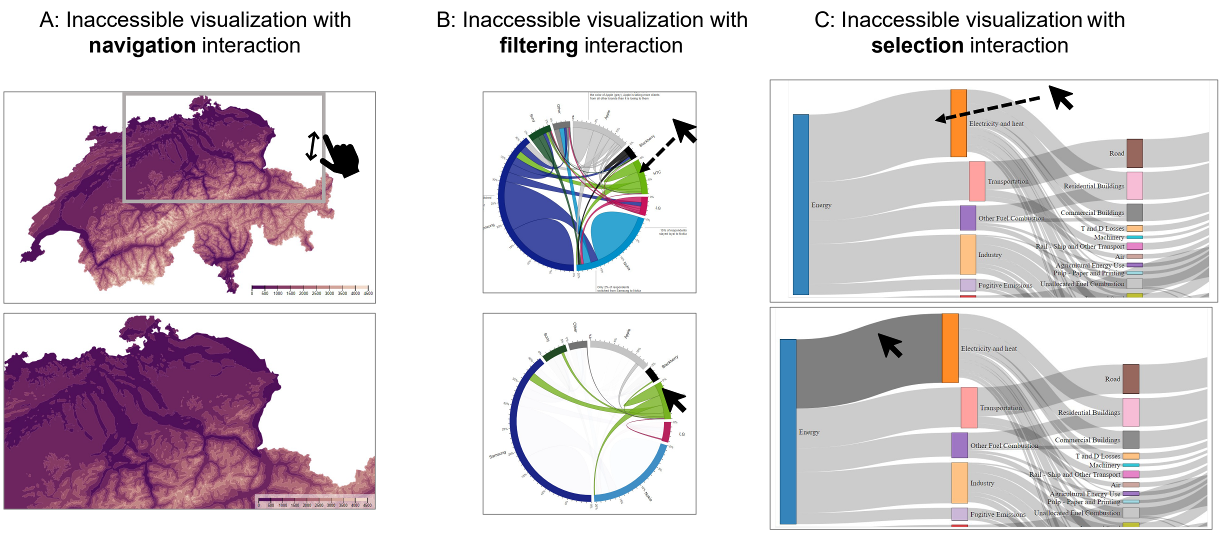

Fig. 1. Three visualizations with navigation (A) [55], filtering (B) [13], and selection (C) [23] interaction mechanisms from our database subset of 375 online D3 visualizations annotated for photosensitive accessibility. Each of the three visualizations in this figure are inaccessible because they are capable of producing flickering sequences that could induce seizures when viewed by people with photosensitive epilepsy.

Abstract—Accessibility guidelines place restrictions on the use of animations and interactivity on webpages to lessen the likelihood of webpages inadvertently producing sequences with flashes, patterns, or color changes that may trigger seizures for individuals with photosensitive epilepsy. Online data visualizations often incorporate elements of animation and interactivity to create a narrative, engage users, or encourage exploration. These design guidelines have been empirically validated by perceptual studies in visualization literature, but the impact of animation and interaction in visualizations on users with photosensitivity, who may experience seizures in response to certain visual stimuli, has not been considered. We systematically gathered and tested 1,132 interactive and animated visualizations for seizure-inducing risk using established methods and found that currently available methods for determining photosensitive risk are not reliable when evaluating interactive visualizations, as risk scores varied significantly based on the individual interacting with the visualization. To address this issue, we introduce a theoretical model defining the degree of control visualization designers have over three determinants of photosensitive risk in potentially seizure-inducing sequences: the size, frequency, and color of flashing content. Using an analysis of 375 visualizations hosted on `bl.ocks.org`, we created a theoretical model of photosensitive risk in visualizations by arranging the photosensitive risk determinants according to the degree of control visualization authors have over whether content exceeds photosensitive accessibility thresholds. We then use this model to propose a new method of testing for photosensitive risk that focuses on elements of visualizations that are subject to greater authorial control – and are therefore more robust to variations in the individual user – producing more reliable risk assessments than existing methods when applied to interactive visualizations. A full copy of this paper and all study materials are available at `https://osf.io/8kzmg/`.

Index Terms—accessibility, photosensitive epilepsy, photosensitivity, interaction, data visualization

1 Introduction

Interaction is a critical component of modern data visualization design. As we visualize larger and more complex datasets with greater computing power, users will often need to apply filters, change views, or drill down in order to effectively gain knowledge [49]. Visualizations that respond rapidly to interaction have been found to facilitate more exploration than those with high latency [36], creating an additional impetus for developers to produce highly responsive visualizations.

• *Laura South and Michelle Borkin are with Northeastern University. E-mail: {south.l, m.borkin}@northeastern.edu*

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxxx/TVCG.201x.xxxxxxx

However, sequences that include sudden changes to color, size, or other aspects of visual appearance are known to be potential triggers for people with photosensitivity. Individuals with photosensitive health conditions, such as photosensitive epilepsy, may experience seizures when exposed to certain light stimuli [25]. The consequences of encountering inaccessible content can be severe for someone with photosensitivity. In addition to seizures, individuals report suffering loss of motor control and neurological symptoms that last for days after exposure [33]. Navigating online spaces while avoiding seizure-inducing content is a challenging task for people with photosensitivity [53]. Social media sites such as Twitter and Reddit often serve GIFs and videos containing dangerous flickering effects without warning. Many individuals do not know they have photosensitive epilepsy until they experience their first light-induced seizure [25], making it even more critical that online content creators ensure that their work is accessible for those with photosensitivity and does not contain dangerous flicker effects.

Prior work has established that data visualizations are capable of producing seizure-inducing sequences as a result of data manipulation [22] and user interaction [52], but to date *no systematic work has been conducted to formalize how data visualizations might be inaccessible to those with photosensitivity*. In this paper we present our work towards addressing this gap. We collected screen recordings of a primary tester interacting with 1,132 online visualizations and evaluated them for seizure-inducing content with a state-of-the-art photosensitive risk detection system for video files called PEAT [19]. When conducting a replicability check with a secondary tester, we found a concerning amount of variation between the PEAT results from the two testers. Visualizations that seemed photosensitive-safe and accessible when interacting with one tester were revealed to be inaccessible and potentially seizure-inducing when used by a different individual with different interaction behaviors.

This variation in photosensitive risk of interactive visualizations between test users is the result of using a system built for analyzing static artifacts with complete authorial control (i.e., GIFs and videos that present the same sequences regardless of the characteristics of the individual doing the testing) to test dynamic artifacts that may depend significantly on the person doing the testing (i.e., interactive visualizations). To effectively evaluate the photosensitive risk of interactive visualizations, *we need a better method that can more reliably determine whether an interactive visualization is photosensitive-safe, regardless of variations among testers*. In this paper we introduce a novel theoretical model defining how flicker can occur in interactive visualizations and how three determinants of photosensitive risk (the size, frequency, and color of a given flicker [26]) may be manipulated in interactive visualizations (Section 4). In particular, the model defines the degree of control visualization designers have over whether a flicker exceeds the size, frequency, and color photosensitive safety thresholds.

Based on this model, we also present a closer look at how visualization designers can create photosensitive-safe color palettes for interactive visualizations and contribute a method for determining photosensitive risk using a JSON (JavaScript Object Notation) specification of visualization interaction techniques. Our method accounts for potential conflicting accessibility needs between those with low vision (who require high contrast colors in visualizations) and those with photosensitivity (who are more likely to be triggered by color changes with high contrast) by simultaneously assessing photosensitive accessibility and contrast accessibility in interactive color palettes.

In summary, this paper makes the following contributions:

- A procedure and script for batch testing screen recordings with state-of-the-art existing risk detection software (PEAT [19]).
- A dataset of interactive visualizations annotated for photosensitive risk.
- A theoretical model of photosensitive risk in interactive data visualizations, building off prior models for photosensitive risk in static media.
- A novel method for testing accessibility of color palettes in interactive visualizations.

We additionally contribute a discussion of strategies for how visualization researchers can make their interactive visualizations more photosensitive accessible. All study materials are available at `https://osf.io/8kzmg/`.

2 Related work

2.1 Accessible data visualization

There has been a recent push to encourage more research on how to make data visualizations accessible for users with a range of abilities [40]. A significant amount of work has focused on accessible visualization for those with visual impairments through textual descriptions [32, 38], data sonification [21], data physicalization [31], and extracting chart features from inaccessible raster images [20]. Other work has examined how color maps in visualizations can be made more accessible to people with color vision deficiency (CVD) [39]. In a reflection on accessible visualization research and its pitfalls, Lundgard et al. emphasized the need for work that uses existing accessibility guidelines to steer and expand research inquiries [37]. Our work builds on prior accessibility guidelines for photosensitivity published in the Web Content Accessibility Guidelines (WCAG 2.0) [17] by extending the established model of photosensitive risk to apply to interactive media that is dependent on the individual user, unlike static GIFs or videos. Wu et al. examined how conventional design guidelines might create obstacles for people with Intellectual and Developmental Disabilities (IDD) when engaging with data visualizations [62]. We take a similar approach in this paper by identifying instances where conventional design guidelines might lead to visualizations that are not accessible for people with photosensitivity, a disability that has not been sufficiently considered in prior work on accessible visualizations.

Flicker and other forms of animation have been examined in data visualization and HCI research as methods for guiding attention [41, 58], facilitating interaction [41, 61], and ameliorating cybersickness [35]. The first mention of flicker in the context of photosensitive accessibility for visualizations occurred when Conti et al. demonstrated that seizure-inducing sequences could be produced by a hypothetical malicious attacker manipulating the data represented in mission-critical visualizations [22]. Similarly, South & Borkin established that seizure-inducing sequences could be produced via interaction alone in web-based visualizations [52]. While both works were valuable in demonstrating the potential for visualizations being inaccessible to people with photosensitivity, neither of the two papers examined in detail how a visualization might be designed to prevent the accidental production of seizure-inducing sequences. We build on this work by contributing an photosensitive accessibility analysis of over 1,000 interactive visualizations, demonstrating the lack of reliability in currently available accessibility testing programs, defining a formal model of photosensitive risk in interactive visualizations, and implementing a novel method for ensuring visualizations are safe for photosensitive users.

2.2 Accessibility for photosensitivity

Research on accessibility for people with photosensitivity has primarily focused on detecting seizure-inducing content in GIFs [53] and videos [4, 7, 18], as both media types have been used to orchestrate malicious attacks against people with photosensitivity on social media (e.g., [3, 16, 33, 46]). The previously mentioned tool PEAT [19] is widely used for checking the safety of videos posted to the internet, while the Harding Flash and Pattern Analyzer [56] is used by broadcast companies to ensure videos shown on television are photosensitive-safe. Several new systems for assessing photosensitive risk have been released in the past ten years, ranging from a rule-based algorithm for real-time detection of seizure-inducing content in video feeds [4] to a machine learning algorithm for removing seizure-inducing content from videos [7]. In this paper, we tackle the issue of photosensitive safety in a new context that has not been considered in prior work: interactive data visualizations.

3 Analysis of interactive and animated visualizations

3.1 Preliminary dataset

In our first step towards understanding the role of photosensitive risk in interactive visualizations, a single coauthor created screen recordings of interactions with 1,132 unique D3 [12] visualizations collected by Hoque & Agrawala with a web crawler [30]. The dataset includes visualizations from a wide range of sources including news websites and digital publications (e.g., New York Times [63], visual essays on The Pudding [47]), and homework assignments[1] and demos of visualization techniques for new developers (e.g., [15]). Each screen recording was then tested for seizure-inducing content [19], an open-source photosensitive risk detection system. PEAT is the industry standard for detecting dangerous flashes in videos and is often used as a benchmark in scientific studies involving the photosensitive risk (e.g., [2, 4, 53]). The procedure used to create all of the screen recordings is included in the Supplemental Material. All recordings were produced by one coauthor

[1] `https://stels07.github.io/05-MapsAndViews/`

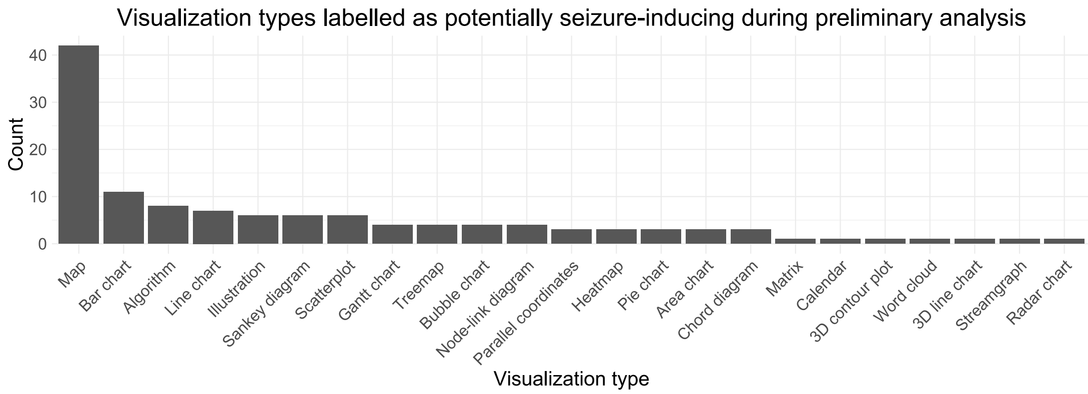

Fig. 2. Frequency of visualization types among visualizations labelled as potentially seizure-inducing during our preliminary analysis (Section 3.1). Our preliminary dataset consisted of screen recordings of 1,132 D3 visualizations tested for seizure-inducing content with PEAT, a state-of-the-art photosensitive risk detection system.

according to a written procedure (included in Supplementary Materials). Recording length varied depending on the complexity of the visualization and the number of interactions implemented by the creator or the length of the animation (minimum duration of two seconds, maximum 69 seconds). This process produced our first contribution: a dataset of 1,132 screen recordings of interactive and animated data visualizations, including 100 recordings (8.8%) that were labelled potentially seizure-inducing by PEAT. The visualizations labelled as potentially seizure-inducing by PEAT spanned a range of visual designs (Figure 2). Maps were most common among dangerous visualizations (42), followed by bar charts (11), animated illustrations of algorithms (8), and line charts (7). All screen recordings and PEAT analysis reports are available at `https://osf.io/8kzmg/`.

3.2 Replicability of testing procedure

To test the replicability of our methodology, we asked a second individual to produce screen recordings of a 10% stratified random sample consisting of 132 visualizations from the main dataset. The secondary tester was a full-time Research Assistant in a data visualization research group with significant course and real-world experience in data visualization comparable to the primary tester. The secondary tester was given the same written procedure to follow that was used by the primary tester when creating recordings (included in Supplementary Material). The secondary tester was not informed whether any of the visualizations were labelled hazardous by PEAT during the primary testing phase. The primary and secondary testers' independent PEAT results had a Cohen's kappa equal to 0.622 (95% confidence interval: $[0.402, 0.842]$), which indicates only moderate agreement [59].

The moderate agreement score across the two testers brought into question the validity of the testing procedure used in our preliminary analysis. This discrepancy demonstrated that *we needed a better way to measure photosensitive risk in interactive visualizations*. We need a method that takes into account the variation in individual behaviors when interacting with the visualizations. In other words, a way of evaluating photosensitive risk based on aspects of a visualization that are not controlled by the tester (i.e., aspects that are subject to authorial control). Existing methods were developed to test media that have total authorial control, such as GIFs and videos, where the content does not depend on the individual viewing it. But the sequences produced by interactive visualizations can vary depending on the person interacting with the system, so testing screen recordings with systems such as PEAT is not a reliable way to tell if a visualization is potentially hazardous or not. In order to more effectively and reliably evaluate the potential photosensitive risk of an interactive visualization, in the next subsection we more closely examine the source of variability between user interaction styles.

3.3 Detailed analysis

After determining that we only had moderate levels of inter-rater agreement when using the testing approach described in Section 3.1, we wanted to investigate the sources of variation that caused certain visualizations to be labelled as potentially seizure-inducing by one rater and not another. To understand more about where the variation between testers originated, we focused on the subset of visualizations hosted on `bl.ocks.org`, henceforth referred to as the **Blocks subset**. We chose to focus on these interactive visualizations as all their source code and data are publicly available. More specifically, regarding sources of variation, we were interested in examining how each visualization handled animated transitions between states in response to interaction, as a user who interacts very quickly with a system might produce a rapid flicker that is more likely to be labelled as seizure-inducing by PEAT than a user who interacts more slowly. For each visualization in the Blocks subset, we recorded what types of transitions were precipitated by user interaction. Our coding was based on Heer & Robertson's taxonomy of transitions [29]. As the taxonomy of transitions does not take into account the transition of an individual element in a non-transitioning visual encoding, e.g., selection and highlighting, we contribute an additional four transition types used to indicate selection that were recurrent in our dataset. We identified 11 **transition types** in our analysis, and present each below with a definition and example for each type:

1. **View transformation/navigation**: A change in viewpoint, similar to a camera moving in space (e.g., panning and zooming in Figure 1A [55] and Figure 5 [11]).
2. **Substrate transformation**: Changes to the spatial substrate in which marks are embedded (e.g., axis rescaling [51]).
3. **Filtering**: Elements are shown or hidden based on data attributes (e.g., Figure 1B [13]).
4. **Ordering/sorting**: Visual elements are rearranged in response to manual input or as a result of automatic sorting based on a data attribute(e.g., [5]).
5. **Timestep**: Temporal changes are applied to data values (e.g., Figure 9 [43]).
6. **Visualization change**: The visualization type or the visual encodings within the same visualization type change in response to user interaction (e.g., [9]).
7. **Data schema change**: The visualized data are changed in some manner, such as adding new data points and attributes or drilling down in a hierarchical data structure (e.g., [6]).
8. **Selection (Fill)**: Changes to a visual element's fill color are used to indicate the user's selection (e.g., Figure 1C [23]).
9. **Selection (Stroke)**: Changes to a visual element's stroke color are used to indicate the user's selection (e.g., [54]).
10. **Selection (Size)**: Changes to a visual element's size are used to indicate the user's selection (e.g., [27]).

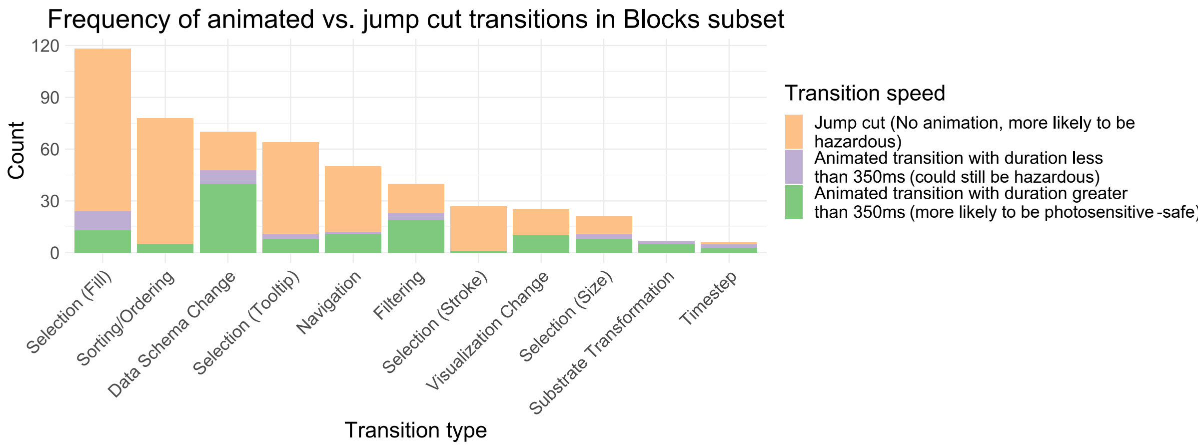

Fig. 3. Visualizations that apply gradual animations with a duration greater than 350ms are less likely to produce flicker effects that can cause seizures when viewed by someone with photosensitivity [28]. We labelled each of the visualizations in the Blocks dataset according to the transition type and whether an animated transition was used to ease between states.

11. **Selection (Tooltip)**: The appearance and/or position of a tooltip is used to indicate the user's selection (e.g., [50]).

With this more detailed labeling of the transition types present in the visualizations included in the Blocks subset, we were able to record whether the transitions were implemented as **"jump cuts"** (i.e., a sudden change between states) or if they were implemented with **gradual animations** easing between state changes. Flashes may be seizure-inducing if they occur more than three times per second [26], assuming additional thresholds for the area and color of the flicker are exceeded (see Section 4.1). Transitions that have a gradual animation lasting more than approximately 350ms are therefore more likely to be accessible for someone with photosensitivity than a visualization that uses abrupt jump cuts for all transitions.

Figure 3 summarizes the proportion of photosensitive-safe animated transitions in our dataset for each transition type observed in our Blocks subset. We recorded 506 transitions in total across 375 visualizations. Selection (Fill) was most common (118 transitions; 23%), followed by Selection (Tooltip) (64 transitions; 12.6%). Stroke and size were less commonly used to indicate a selection, occurring only 27 (5.3%) and 21 (4.2%) times. Following Selection (Fill), Sorting/Ordering and Data Schema Change were the second and third most common transition types, appearing 78 (15.4%) and 70 (13.8%) times.

Approximately two-thirds of transitions observed in the Blocks subset used jump cut transitions, meaning there was no animation easing between states of the transition (349 out of 506 total transitions; 68.8%). Jump cut transitions were most common among Sorting/Ordering (73 out of 78 transitions; 94.5%), Selection (Stroke) (26 out of 27 transitions; 96.3%), Navigation (38 out of 50 transitions; 76%), and Selection (Tooltip) (53 out of 64 transitions; 82.8%) transition types. Jump cut transitions are not necessarily inaccessible for those with photosensitivity – photosensitive risk also depends on the size and color of a given flicker (Section 4) – but they can help to limit how rapidly a flicker effect repeats within an interactive visualization. The lack of gradual animated transitions in the Blocks subset may begin to explain the variation among testers we found in Section 3.1.

Even among visualizations that did incorporate gradual animated transitions to ease between transition states, 34 transitions were implemented with a duration that was too short to effectively prevent a dangerous flicker from occurring. As we will discuss in Section 4.3, a gradual animated transition must last at least 350ms to ensure that the photosensitive safety threshold for frequency is not exceeded (i.e., the flash does not occur more than three times per second). This means that a particularly vigorous user interacting quickly with one of these visualizations could still hypothetically generate seizure-inducing sequences, despite their use of gradual animated transitions.

Our detailed analysis of the Blocks subset reveal that although gradual animated transitions might remove some of the user variability that we observed when testing screen recordings of interactive visualization with PEAT, most visualizations did not include such transitions, or included animated transitions that were too short to ensure photosensitive accessibility. By analyzing the Blocks subset and looking for sources of user variability, we developed a theory of authorial control over photosensitive risk in visualizations. This theory, which is described in detail in Section 4, defines the level of control a visualization creator holds over each of the three determinants of photosensitive risk (size, frequency, and colors of a flicker). We use examples from the Blocks subset to demonstrate how visualization creators have little control over the size of a flicker in an interactive visualization, but have almost complete control over the colors included in a flicker.

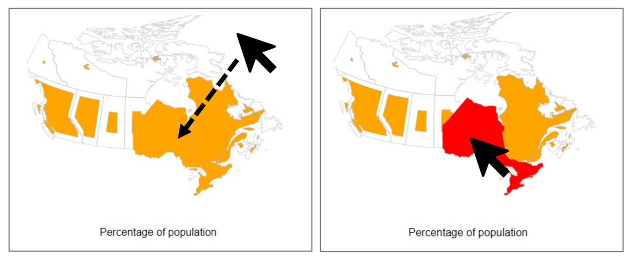

Fig. 4. A map [42] with the Selection (Fill) transition demonstrates how internal element flicker occurs when a visual element changes color in response to user interaction.

4 Model of Interactive Photosensitive Risk

In this section, we describe the conventional model of photosensitive risk, which defines three thresholds (size, frequency, and color) that a flicker must exceed to be considered potentially seizure-inducing [28]. This model has been successfully used to detect seizure-inducing content in static media formats, such as videos and GIFs (e.g., [4, 19]), but has not previously been applied to interactive media formats. Using observations from our dataset of interactive visualizations (Section 3.3, we extend this model to account for variation among users inherent in interactive data visualizations. We begin by defining two ways that flicker can appear in interactive visualizations and continue by introducing the three determinants of photosensitive risk and examining their behavior in the context of interactive data visualizations.

4.1 Flicker types

A **flash** is a pair of opposing changes in relative luminance (i.e., light-dark or dark-light). We will use the term **flicker** to refer to a series of

Flicker attribute	Photosensitive safety threshold	Mitigation techniques
Size (Section 4.2)	Occupies less than 10% of the central visual field (approximately 200-300px square on computer screen or 50-70px square on mobile device [24])	N/A
Frequency (Section 4.3)	Occurs at a rate of less than 3 flashes per second	Gradual animated transitions
Color (Section 4.4)	$L_1 - L_2 < 0.1 \times L_1$ or $L_2 > 0.8$, where L_1 is the relative luminance of the brighter color incorporated in the flash and L_2 is the relative luminance of the darker color	Use colors with similar levels of saturation and luminance in flickering elements

Table 1. Summary of three attributes of a flicker effect that determine whether it is considered a photosensitive hazard [26]. All three thresholds (size, frequency, and color) must be exceeded for a flicker to be considered seizure-inducing.

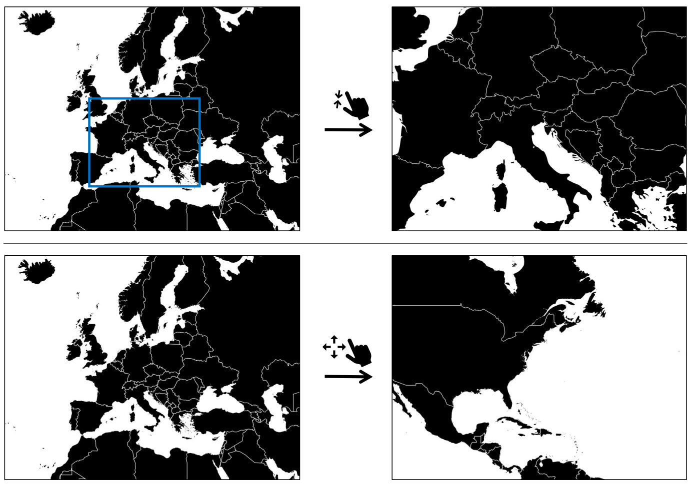

Fig. 5. An example of adjacent element flicker found in an interactive D3 map [11]. Although individual elements do not change color on interaction, a flicker effect occurs as the black landmasses and white oceans alternate covering the same screen area when the user repeatedly pans (bottom) and zooms (top).

flashes. A flicker can cause a seizure if it occupies a large enough area, occurs at the right frequency, and incorporates a sharp enough change in relative luminance. Flicker can occur in interactive visualizations in two ways: First, a visualization might produce **internal element flicker** if the color of visual elements change in response to user interaction. A common example of a singular flash that could produce a flicker when repeated several times is when a visual element changes color to show a user their current selection in a chart (Figure 4). Second, a visualization can produce **adjacent element flicker** if elements may be rearranged or resized within a visualization. Even if the color of the elements is not directly changed by user interaction, a flicker effect can be formed as a result of rapid movement of multiple elements against a constant background. Interactive maps are a common source of adjacent element flicker (e.g., Figure 5).

Accessibility guidelines (e.g., WCAG Success Criterion 2.3.1 [17]) have used EEG-based empirical studies measuring brain response to establish thresholds for three determinants that control the seizure-inducing potential of a flicker: size, frequency, and color [26] (Table 1). A flicker that exceeds all three risk factor thresholds is likely to cause a photosensitive response when viewed by someone with photosensitivity. These thresholds are not dependent on individual screen settings, such as brightness and contrast. Content with flashes may be less likely to cause seizures if viewed on a screen with low brightness, but a designer looking to make accessible content must ensure that their work is accessible even in a hypothetical scenario where it is viewed at maximum brightness. In the remaining sections we will summarize the three determinants included in the traditional model of photosensitive risk for static media (e.g., GIFs or videos) and discuss how they may be applied to interactive data visualizations.

4.2 Size

A flicker may be dangerous if the total area of concurrent flashes occupies a solid visual angle of ≥ 0.006 steradians, which equates to approximately 10% of the central visual field or 25% of the area of a television screen at a standard viewing distance [28]. Most photosensitive risk thresholds are defined in the context of viewing content on a television screen because they predate mobile device usage, but the formulas can be extended to account for other screen types. Assuming standard viewing distances on desktops and mobile devices, flicker may be dangerous if it occupies more than a 200-300px square on a desktop monitor or more than 50-70px square on a mobile device [24].

In terms of interactive data visualizations, this means that elements that change colors in a visualization are more likely to be dangerous if they occupy a large area (e.g., areas in an area chart) rather than a small area (e.g., points on a scatterplot). Visualization creators can control the initial size of visual elements and their relative size within the visualization. They can also place limits on zoom behavior to restrict how large an element can become (e.g., D3's scaleExtent) or choose not to implement zooming at all in the visualization. But once the user is interacting with the visualization they can accidentally or intentionally zoom in using the web browser, potentially bypassing any zoom limits set by the creator.

Because all three thresholds must be exceeded for content to be considered a photosensitive hazard, a visualization designer might choose to restrict color changes to smaller elements in a visualization. Such a visualization might pass a traditional photosensitive risk test with a system like PEAT, where the designer creates a screen recording of themselves interacting with the system as they intended it to be used. But if a different person using the visualization is not accustomed to the interaction techniques implemented in it, they might accidentally use the browser to increase the size of visual elements. Depending on whether the frequency and color thresholds (Sections 4.3 and 4.4) are exceeded, a visualization that was safe in the hands of its creator might now pose a risk to photosensitive users. In other words, *a visualization that is safe only because flickering visual elements are small is not reliably safe in all scenarios*. The size threshold for photosensitive risk is **not robust** for ensuring the safety of internal element flicker or adjacent element flicker in interactive data visualizations.

4.3 Frequency

A flicker may be hazardous if it includes more than 3 flashes (i.e., pairs of opposing changes in relative luminance) in any one-second period. Any flicker that occurs at a rate greater than 3 times per second may be considered hazardous, depending on whether the other two thresholds are exceeded. Different users will interact with visualizations at different speeds, which will in turn produce internal element flicker and adjacent element flicker at different frequencies, but a degree of authorial control over flicker frequency can be found in some scenarios. Visualization creators can exert authorial control over the maximum frequency of a given flicker by attaching gradual transitions with a duration greater than 350ms to user interaction events (e.g., D3 transitions[2]), thereby preventing the flicker from occuring more than three times per second. Transitions can be used to ease adjacent element flicker by attaching transitions directly to pan and zoom transformations of the visual representation. Gradual transitions are not appropriate in all scenarios; interaction lags may decrease user engagement [36] or impact usability in a system that requires many successive interactions.

In summary, the frequency threshold applied to interactive data visualizations can be **robust** for ensuring the safety of a flicker if gradual transitions are used to ease animations between states, but delayed transitions could introduce additional usability problems in some applications by preventing quick interactions.

4.4 Color

According to accessibility guidelines, the photosensitive risk present in a given flicker is in part determined by the two colors (one lighter and one darker) that form the flash. More specifically, photosensitive risk is determined by the difference in relative luminance between the brighter color (L_1) and the darker color (L_2). Relative luminance is a normalized value between 0 and 1 representing the relative brightness of any point in a colorspace [1]. If $L_1 - L_2 > 0.1 \times L_1$ and $L_2 < 0.8$, the general flash may be hazardous. Many people with photosensitivity are particularly sensitive to the color red [8], so red flashes are considered separately from general flashes. Accessibility guidelines do not currently give an explicit definition of safety thresholds for red flashes, so we focus on general flashes in this work.

In most cases, a visualization creator has complete authorial control over the colors present in any flicker that arises from their visualization. Unlike the size and frequency of a flicker, a user cannot easily make changes to the colors shown in a visualization that produce a flicker effect. This makes color the most **robust** threshold in terms of controlling photosensitive risk in interactive visualizations. A visualization that has photosensitive-safe colors for all possible internal element flickers and adjacent element flickers will be robustly safe regardless of how an individual user chooses to interact with the visualization.

4.5 Model Summary & Discussion

Three characteristics determine whether a flicker is considered a photosensitive risk (i.e., capable of causing a seizure): size, frequency, and color. Each characteristic has a specific safety threshold that is used to determine if content is hazardous or not. All three thresholds must be exceeded for a sequence to be considered inaccessible for someone with photosensitivity. In the preceding section, we have defined a model connecting each of the three determinants of photosensitive risk to the domain of interactive visualizations. In particular, we demonstrated that the size threshold can be easily bypassed by a user, accidentally or intentionally, by zooming in on elements in the browser (Section 4.2). The frequency and color thresholds, in contrast, can be robustly controlled by visualization creators to ensure that visualizations are not capable of producing seizure-inducing flickers. The frequency at which a flicker occurs in a visualization can be limited to stay below the safety threshold by applying a gradual transition with a duration of at least 350ms (Section 4.3). This can be easily implemented in D3 visualizations using the d3-transitions functionality[3]. The color safety threshold is also robust to variations among individual users, as users cannot easily change the colors in a visualization beyond what was originally intended by the creator. However, no tool currently exists to automatically assess the photosensitive accessibility of the colors used in a given visualization. In the next section, we introduce our final contribution: a method for analyzing photosensitive risk based on the colors and interaction techniques used in a visualization.

5 PHOTOSENSITIVE-SAFE COLOR PALETTES

In the previous section, we discussed three characteristics of a flicker (size, frequency, and color) that determine whether or not it is considered seizure-inducing. Our model of photosensitive risk in interactive

[2]https://github.com/d3/d3-transition
[3]https://github.com/d3/d3-transition

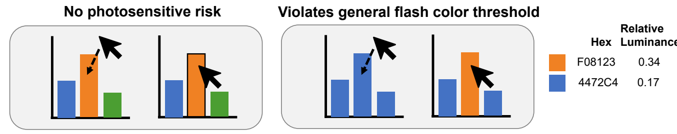

Fig. 6. An example demonstrating how the same color palette (in this case, D3's Category10 colormap) can be used to create visualizations that are accessible (left) and not accessible (right) for those with photosensitivity.

visualizations defined the degree of control that visualization creators have over each determinant of photosensitive risk and identified color as an aspect of photosensitive risk that is almost entirely controlled by the visualization creator. The size of a flickering element in an online visualization cannot be controlled by the designer because users can use the browser to change the size of the visualization, intentionally or otherwise. Similarly, if a visualization does not incorporate gradual animated transitions, designers have no control over how quickly the user interacts with the system and, by extension, little control over the frequency of any flickers present in the visualization. Color, on the other hand, is entirely controlled by the visualization designers and cannot be changed by the user in most instances.

A visualization with photosensitive-safe colors is robustly safe regardless of individual interaction behavior, but determining whether the colors in a visualization are capable of producing dangerous flicker is a nontrivial problem. The same colors used in various ways can produce visualizations with vastly different levels of photosensitive risk. As demonstrated in Figure 6, two bar charts that both use D3's categoryScheme10 categorical color palette in different ways result in one safe and one hazardous interactive visualization. The first uses the categorical color palette to fill each bar and a black stroke outline provides visual feedback to the user about their current selection. The second uses the first color in the categorical palette as the initial fill color and changes the fill to an orange color to indicate the user's selection. The first bar chart has photosensitive-safe colors because the area that changes color is very small, while the second bar chart could be hazardous due to the change in fill between blue and orange on selection. In this section, in order to enable visualization designers to evaluate the accessibility of their creations, we contribute and describe the implementation of a Python script for testing photosensitive safety of color palettes in interactive visualizations. The script and examples of its use are both available in Supplemental Material or at `https://osf.io/8kzmg/`.

5.1 Assigning color roles

To determine the photosensitive safety of a given color scheme, we need to know the colors in a visualization *and* their roles in the interaction mechanisms of the visualization. This allows us to locate where internal element flicker and adjacent element flicker might occur and make changes to color schemes to ensure that the flickers do not exceed the color thresholds described in Section 4.4. We can define all colors capable of contributing to an internal element or adjacent element flicker in a visualization using the JSON (JavaScript Object Notation) specification demonstrated in Figure 7B. Each element in the specification must have a unique ID and an `initial-fill` (i.e., a hex code representing the color of the element prior to any user interaction). The visualization specification may include a `background` element (indicated by the ID attribute) to specify a background color other than the customary white (`#FFFFFF`). Elements may also have the following optional attributes to capture more specific interaction behaviors:

- `focus-fill`: One or multiple colors other than initial fill that the element can have in response to user interaction.
- `focus-opacity`: An alternate method of specifying a focus fill, in which the resulting focus fill is calculated by blending the initial fill and background with the provided opacity value.
- `initial-stroke`: If the visual elements have a stroke that differs from the initial fill, it can be specified here.

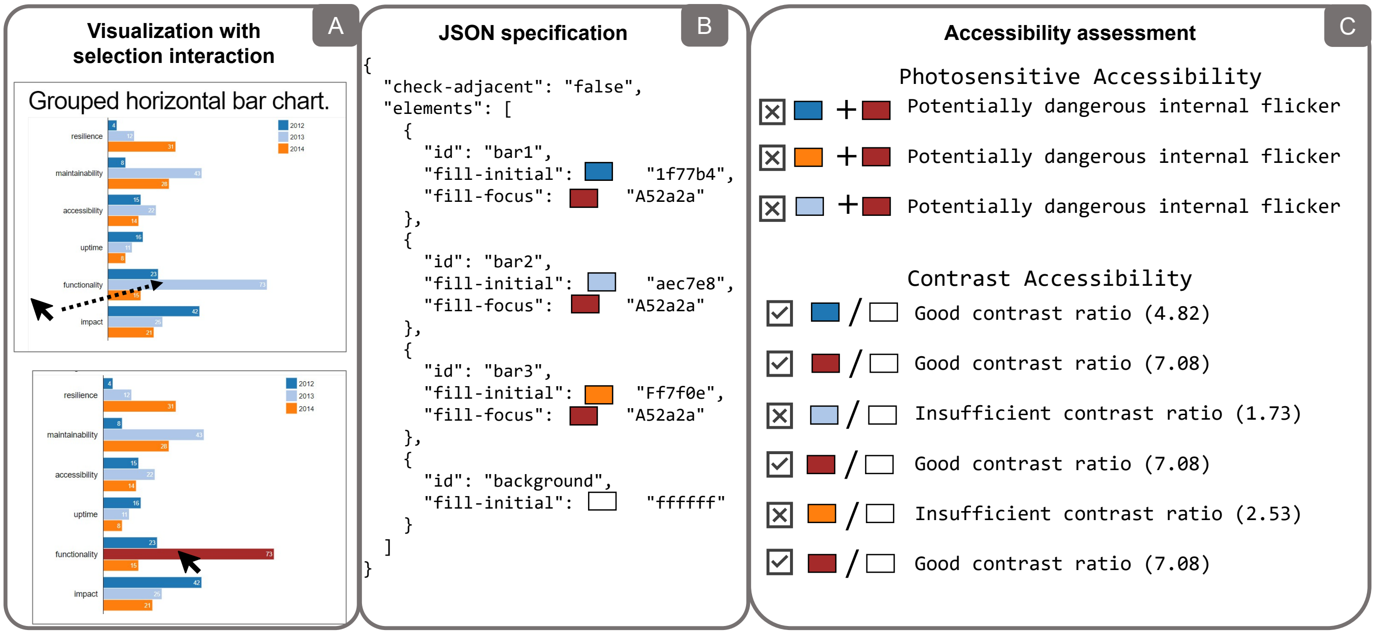

Fig. 7. An interactive grouped bar chart [57] (A) capable of producing color-change sequences that may trigger seizures for individuals with photosensitivity. We propose a method of measuring photosensitive risk that analyzes the colors used in a visualization specified in JSON format (B) and summarizes the visualization's accessibility for those with photosensitivity while simultaneously ensuring that visual elements have sufficient contrast for users with low vision (C).

- `focus-stroke`: If the stroke of an element changes color on focus, it can be specified here.

Next we will discuss how these attributes are used to calculate photosensitive risk potential for color palettes in interactive visualizations.

5.2 Internal element flicker

Elements with internal element flicker change color in response to user interaction (Section 4.1). This behavior can be represented in the specification by defining one or multiple `focus-fill` attributes within the current element. All pairs of `initial-fill` and `focus-fill` colors within the same element will be checked for photosensitive risk to ensure they will not produce a seizure-inducing internal element flicker. While some creators explicitly name a secondary color for an element to change to on user interaction, many visualizations implicitly change the fill color by varying the opacity attribute of the element. This behavior can be replicated in our specification by defining the `focus-opacity` attribute, which should be a value between 0 and 1 and will be used to blend the element's `initial-fill` and the `background` fill to calculate the element's ultimate `focus-fill`. Once all elements with internal flicker have both an `initial-fill` and `focus-fill` colors, we can check the photosensitive risk of each color combination using the color threshold from Section 4.4.

5.3 Adjacent element flicker

Adjacent element flicker occurs when elements can be moved around on the screen quickly enough that a flicker effect appears as the elements and the background successively occupy the same area of the screen (Section 4.1). Not all visualizations are capable of producing an adjacent element flicker. Visualizations that do not allow the user to rearrange, reshape, or otherwise alter the position of visual elements (manually or otherwise) do not need to be checked for adjacent element flicker. Therefore we include an attribute (`check-adjacent`) to indicate whether a visualization is capable of producing adjacent element flicker. If the visualization needs to be checked for adjacent element flicker, we need to examine the risk potential of all possible color combinations. This means we need to look at all initial and focus fills and test them pairwise against the color threshold. We also need to check all initial and focus fills against the background color.

5.4 Additional accessibility constraints

It is important that we do not introduce new accessibility issues into visualization designs in our attempts to remove or mitigate dangerous flicker. There is a potential for contradicting accessibility requirements between people with low vision and people with photosensitivity. Individuals with low vision are more likely to be able to read a visualization with highly saturated colors and a lot of contrast, while people with photosensitivity are more likely to be triggered by visualizations with bright colors if an interactive flicker effect is present in the visualization. Figure 8A shows a map with maximum contrast (black landmasses against a white background). Because this visualization was not implemented with gradual animations, a dangerous flicker can occur when the user pans and zooms to navigate the map. Figure 8B shows the same visualization edited to have photosensitive-safe colors. In the process of making the visualization accessible for those with photosensitivity, we have made the contrast so low that it is inaccessible for those with low vision. Figure 8C keeps the low-contrast colors from B but adds a high contrast stroke outlining the landmasses and thereby satisfying the two conflicting accessibility needs.

Success Criterion 1.4.11 in the WCAG 2.0 recommends using colors with a contrast ratio greater than 3:1 for graphical objects conveying information on a webpage. The contrast ratio of two colors ($(L_1 + 0.05)/(L_2 + 0.05)$, where L_1 is the relative luminance of the lighter color and L_2 is the relative luminance of the darker color) represents the change in relative luminance between two colors. We can calculate the contrast ratio between all possible pairs of adjacent colors defined in our specification (Section 5.1). Consequently, the script can be used to check **contrast accessibility** even for visualizations that do not have either flicker type. Figure 9) demonstrates how a visualization without internal element or adjacent element flicker can still be tested for contrast accessibility using our color analysis method.

5.5 Case studies

To demonstrate the benefits of our photosensitive-safe color analysis method, we created JSON specifications for several visualizations from the Blocks subset (Section 3.3) and tested them for photosensitive risk. JSON specifications for all tested visualizations and accessibility results are included in our Supplemental Material and at

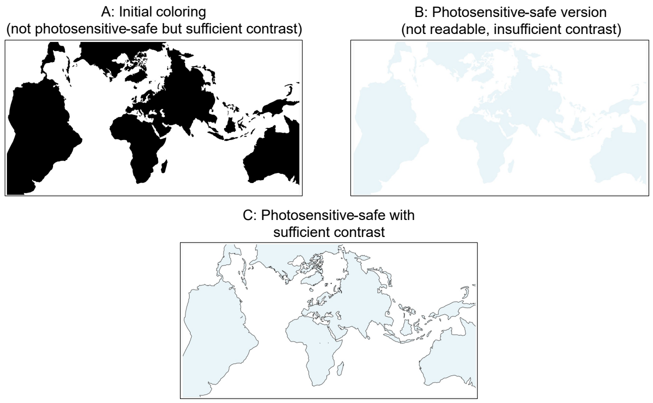

Fig. 8. An interactive map with high contrast colors (A [10]) is accessible for users with low vision, but can be inaccessible for people with photosensitivity if incorporated into elements that produce a flicker effect. Similarly, photosensitive-safe colors may not be accessible for people with low vision (B). Using a high contrast stroke color to separate elements (C) can help balance conflicting accessibility constraints.

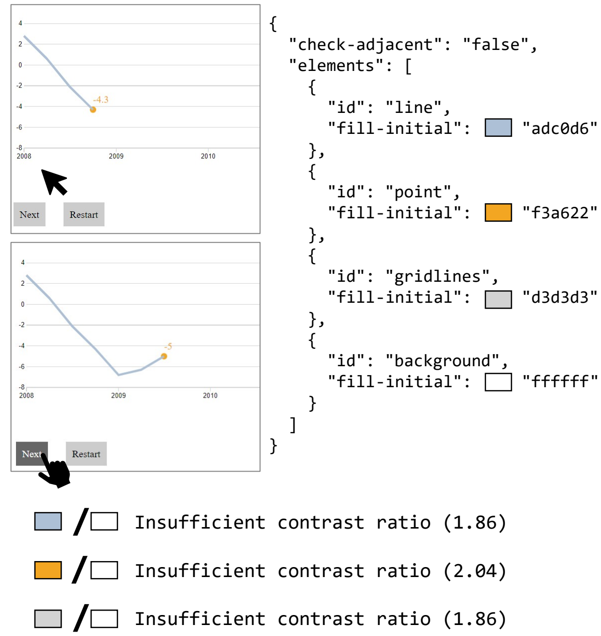

Fig. 9. An example of our color analysis method applied to a visualization [43] with no potential for internal element or adjacent element flicker, as the user can only click a button to advance the animation. In this situation, the method simply assesses the contrast accessibility of the visualization.

https://osf.io/8kzmg/. Table 2 summarizes our ten case studies, each of which appears as examples in the figures of this paper. Eight of the ten case studies were included in the preliminary analysis. The remaining two case studies (8 and 9) are edited versions of case study 7, adjusted to be accessible for photosensitivity (8) and accessible for both photosensitivity and low vision (9). Only one case study (10) was implemented with a sufficiently gradual transition to ensure that the frequency safety threshold would not be exceeded (Section 4.3).

The case studies that were edited to improve accessibility (8 and 9) were both found to be photosensitive-safe. Case study 10 was implemented with gradual transitions and labelled safe by both PEAT and the novel color analysis method. Of the eight case studies that were tested during preliminary analysis, only two were found to be inaccessible using the screen recording procedure described in Section 3.1. When analyzed with the color analysis method, case studies 1-7 were all found to have color combinations capable of producing seizure-inducing flicker if not used in combination with gradual animated transitions. The novel color analysis method allowed us to identify five potentially inaccessible visualizations that were mistakenly labelled as accessible when analyzed using existing systems for detecting photosensitive risk (i.e., screen recording testing procedure described in Section 3.1). By focusing on analyzing the color and frequency of flickers present in these visualizations, we are able to obtain reliable estimates of photosensitive risk that are not dependent on the individual running the test, thereby addressing the issue of reliability identified in Section 3.2.

6 Discussion

6.1 Photosensitive accessibility guidelines

To empower visualization designers to create more accessible visualizations, we present the following guidelines for limiting potential for photosensitive risk in interactive visualizations.

1. **Animated transitions:** Inserting gradual animations to ease between state changes during transitions is a powerful way of ensuring that visualizations are accessible for people with photosensitivity. Animated transitions are supported natively by D3 and can be added into existing D3 visualizations with only a few lines of extra code. Transitions can be used to limit the frequency at which both internal element flicker and adjacent element flicker occur (Section 4.1) in most D3 visualizations. Crucially, animated transitions must have a duration greater than 350ms to ensure that the frequency safety threshold is not exceeded. A D3 transition with no specified duration currently defaults to a 250ms duration[4], which could still produce seizure-inducing flicker that exceeds the three flashes per second threshold, so designers must carefully specify a transition duration greater than 350ms to ensure accessibility.

2. **Photosensitive-safe color palettes:** Animated transitions may not be a feasible solution for visualizations that require rapid response to user interaction or are implemented with technologies that do not have animated transitions easily available. In these scenarios, designers can select color palettes that do not exceed the color threshold for flicker using the method described in Section 5. Designers should carefully consider whether internal element flicker, adjacent element flicker, or both, are present in their visualizations and check the photosensitive risk potential of colors involved in flickers accordingly. Selecting colors that have sufficiently high contrast to be discernible yet do not present a photosensitive risk can be challenging for adjacent element flicker, as demonstrated in Figure 8B. A high contrast stroke outlining the separation between elements can be helpful in balancing the requirements of low-vision and photosensitive accessibility in visualizations (Figure 8C).

6.2 Limitations & future work

While we are confident that this paper contributes an accurate and comprehensive analysis of photosensitive risk in the interactive visualizations included in our survey, we acknowledge several potential limitations of our approach. First, the visualizations included in our full dataset were collected via web crawler by Hoque & Agrawala [30] and may not be fully representative of the entire design space of interactive visualizations. By definition, our dataset only includes visualizations implemented with D3 and therefore our results may not generalize to visualizations created with other technologies, such as HTML5's Canvas or Vega-Lite [48]. Furthermore, the visualizations included in our detailed analysis subset (Section 3.3) are all hosted on bl.ocks.org

[4] https://github.com/d3/d3-transition

ID	Location in paper	Visualization type	Transition type(s)	Transition animation style	Preliminary analysis result (PEAT)	Novel color analysis result
1 (I-123)	Figure 1A [55]	Map	Navigation	No animation	Dangerous	Dangerous
2 (I-178)	Figure 1B [13]	Chord diagram	Filtering	Duration <350ms	Safe	Dangerous
3 (I-77)	Figure 1C [23]	Sankey diagram	Selection (Fill)	No animation	Safe	Dangerous
4 (I-243)	Figure 4 [42]	Map	Selection (Fill)	No animation	Safe	Dangerous
5 (I-479)	Figure 5 [11]	Map	Navigation	No animation	Dangerous	Dangerous
6 (I-110)	Figure 7 [57]	Bar chart	Selection (Fill)	No animation	Safe	Dangerous
7 (I-443)	Figure 8A [10]	Map	Navigation	No animation	Safe	Dangerous
8 (N/A)	Figure 8B [10]	Map	Navigation	No animation	N/A	Safe
9 (N/A)	Figure 8C [10]	Map	Navigation	No animation	N/A	Safe
10 (I-30)	Figure 9 [43]	Line chart	Timestep	Duration >350ms	Safe	Safe

Table 2. We applied the color analysis method described in Section 5 to ten case studies from our Blocks subset and identified five visualizations that were originally mislabelled as accessible in the preliminary analysis. JSON specifications and accessibility results for each visualization, along with all data collected during preliminary and detailed 3) are available at `https://osf.io/8kzmg/`. Full metadata for the preliminary dataset and the Blocks subset are additionally available at `https://airtable.com/shrRbuna0dz0a3hl9`. Blocks dataset IDs are listed in parentheses.

and may be less complex than visualizations hosted on other websites. Nevertheless, there is intellectual merit in demonstrating how the methods we have developed can be applied to relatively simple visualization designs, particularly as many of these straightforward designs are commonly used in visualization practice.

Our work has centered around identifying and mitigating seizure-inducing flicker and does not discuss the potential for seizure-inducing patterns in interactive or static data visualizations. Pattern sensitivity (i.e., experiencing seizures in response to pattern stimuli such as high-frequency stripes in addition to light stimuli) is a known phenomenon among those with photosensitivity and adverse effects can occur in response to even static images, depending on the characteristics of any repeated patterns present in the image [60]. Furthermore, this work relies heavily on the three-determinant model of photosensitive risk defined by Harding et al. in 2005 [28] (Section 4), which is by necessity a simplification of the actual physiological processes underlying a photosensitive response. Recent work has focused on more accurately estimating the probability of visual stress using continuous variables [34, 44, 45], rather than the discrete variables used in the current photosensitive risk model [28]. A similar approach may yield more accurate models of photosensitive risk in the future. There is potential for valuable future work examining how more sophisticated models of flash and pattern sensitivity might relate to the design of accessible static and interactive data visualizations.

6.3 Conflicting design needs

We discuss in Section 5.4 the potential for conflicting requirements between designing visualizations that are accessible for people with low vision and people with photosensitivity, and contribute a method for mitigating photosensitive risk in color palettes without sacrificing high contrast for low vision legibility. Our work also reveals several areas where established visualization design idioms might go against the needs of people with photosensitivity and may inadvertently lead to the proliferation of inaccessible visualizations. First, empirical research has shown that systems that respond to user interaction without delays or lags are more effective [36]. Our work indicates that there could be accessibility benefits in having a system that is slightly slower to respond to user input. Gradual animated transitions with a duration greater than 350ms can help ensure that any flicker related to user interaction in a visualization will not exceed the frequency threshold (Section 4.3) and will be less likely to trigger seizures. This will additionally support better graphical perception of change during a visualization transition [29].

Additionally, perceptual research shows that saturation and luminance are powerful channels for encoding sequential information or creating areas that are easily distinguishable from each other [14]. Color palettes that vary levels of saturation and luminance are generally more accessible for people with color-vision deficiency (CVD), i.e., color blindness, as differences in hue are more difficult to detect for those with CVD. At the same time, our work demonstrates that color maps with large differences in saturation and luminance may be inaccessible for those with photosensitivity if the difference in luminance is sufficiently large that the color threshold described in Section 4.4 is exceeded. We mention this not to argue against using color maps that vary saturation and luminance, but rather to bring attention to additional accessibility constraints which designers should keep in mind when designing interactive visualizations to ensure they can safely be used by as many individuals as possible.

6.4 Call to action

As we detail in Section 2.1, accessibility research within the data visualization community has historically focused on building accessible charts for people with CVD (i.e., color blindness), or vision impairments. The number of potential users with photosensitive epilepsy or other forms of photosensitivity may be small, but the effects of encountering an inaccessible visualization with seizure-inducing flicker can be debilitating for an affected user. We hope that our work demonstrates the need for inquiry into the current state of visualization accessibility for a broader range of users. There are many levels of ability that have not been considered in prior research on visualization accessibility, particularly cognitive and neurological disabilities. The potential for intellectual merit and positive impact on the lives of users is vast and should not be underestimated. We hope that our work can serve as a template for future research related to interpreting existing accessibility guidelines for visualization design, identifying gaps in currently available tools for measuring accessibility, and developing new methods to make it easier for the data visualization community to align itself with existing accessibility guidelines.

7 Conclusion

We have presented the first systematic analysis of photosensitive accessibility in the context of interactive data visualizations. We have shown that prior systems built to assess photosensitive risk in videos and screen recordings are unreliable when applied to interactive visualizations and may produce varied risk assessments when tested with different users. We contributed a theoretical model for conceptualizing photosensitive risk in interactive visualizations and a novel method for measuring photosensitive risk based on the implementation details of a visualization. We hope that our work illustrates the potential for future research into understanding how visualizations can be made more accessible for people of all abilities and backgrounds.

Acknowledgments

We would like to thank members of the Northeastern Visualization Lab for their helpful feedback and support. This research is supported in part by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE1451070, and the Khoury College of Computer Sciences, Northeastern University.

References

[1] WCAG definition of relative luminance. *WCAG Working Group*.

[2] P. Acosta-Vargas, L. Salvador-Ullauri, J. L. Pérez-Medina, and Y. Rybarczyk. Accessibility evaluation of multimedia resources in selected latin america universities. In *2019 Sixth International Conference on eDemocracy & eGovernment (ICEDEG)*, pp. 249–255. IEEE, 2019.
[3] J. Aker. Epilepsy Foundation files criminal complaint and requests investigation in response to attacks on Twitter feed. *Epilepsy Foundation*, Dec. 2019.
[4] M. A. Alzubaidi, M. Otoom, and A.-K. Al-Tamimi. Parallel scheme for real-time detection of photosensitive seizures. *Computers in Biology and Medicine*, 70:139–147, 2016.
[5] N. Atkinson. Bi-directional hierarchical sankey diagram, Mar. 2020. http://bl.ocks.org/Neilos/584b9a5d44d5fe00f779/.
[6] M. Baiges. Pie chart drill down test, Sept. 2019. http://bl.ocks.org/marcbc/3281521.
[7] A. Barbu, D. Banda, and B. Katz. Deep video-to-video transformations for accessibility with an application to photosensitivity. *Pattern Recognition Letters*, 2019.
[8] J. Bhattacharya. Seeing red: Colour modulation and photosensitive epilepsy. *Epilepsy Professional*, pp. 26–29, 2010.
[9] M. Bostock. Stacked-to-multiples, June 2017. https://bl.ocks.org/mbostock/4679202/.
[10] M. Bostock. Interactive stereographic, Nov. 2018. https://bl.ocks.org/mbostock/3763057/.
[11] M. Bostock. Constrained zoom, Apr. 2020. https://bl.ocks.org/mbostock/4987520/.
[12] M. Bostock, V. Ogievetsky, and J. Heer. D^3 data-driven documents. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2301–2309, 2011.
[13] N. Bremer. Chord diagram - phone brand switching, June 2016. http://bl.ocks.org/nbremer/6540350.
[14] C. A. Brewer et al. Color use guidelines for data representation. In *Proceedings of the Section on Statistical Graphics, American Statistical Association*, pp. 55–60, 1999.
[15] P. Buffa. 4.0 color scales - sequential. *bl.ocks.org*, Feb. 2020.
[16] M. Busby. Malicious tweets targeting epilepsy charity trigger seizures. *The Guardian*, May 2020.
[17] B. Caldwell, M. Cooper, L. G. Reid, and G. Vanderheiden. Web Content Accessibility Guidelines (WCAG) 2.0. *WWW Consortium (W3C)*, 2008.
[18] L. Carreira, N. Rodrigues, B. Roque, and M. P. Queluz. Automatic detection of flashing video content. In *2015 Seventh International Workshop on Quality of Multimedia Experience (QoMEX)*, pp. 1–6. IEEE, 2015.
[19] T. R. Center. Trace center photosensitive epilepsy analysis tool (peat), 2021.
[20] J. Choi, S. Jung, D. G. Park, J. Choo, and N. Elmqvist. Visualizing for the non-visual: Enabling the visually impaired to use visualization. In *Computer Graphics Forum*, vol. 38, pp. 249–260. Wiley Online Library, 2019.
[21] P. Chundury, B. Patnaik, Y. Reyazuddin, C. Tang, J. Lazar, and N. Elmqvist. Towards understanding sensory substitution for accessible visualization: An interview study. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):1084–1094, 2021.
[22] G. Conti, M. Ahamad, and J. Stasko. Attacking information visualization system usability: overloading and deceiving the human. In *Proceedings of the 2005 Symposium on Usable Privacy and Security*, pp. 89–100, 2005.
[23] D3noob. Sankey diagram with horizontal and vertical node movement, Nov. 2019. http://bl.ocks.org/d3noob/5028304/.
[24] M. W. Docs. Web accessibility: Understanding colors and luminance, Feb. 2022.
[25] G. Erba. Shedding light on photosensitivity, one of epilepsy's most complex conditions. *Epilepsy Foundation*, 2006.
[26] R. S. Fisher, G. Harding, G. Erba, G. L. Barkley, and A. Wilkins. Photic- and pattern-induced seizures: a review for the Epilepsy Foundation of America Working Group. *Epilepsia*, 46(9):1426–1441, 2005.
[27] J. Freels. d3js: Scatterplot using csv data, June 2018. http://bl.ocks.org/jfreels/6816504/.
[28] G. Harding, A. J. Wilkins, G. Erba, G. L. Barkley, and R. S. Fisher. Photic-and pattern-induced seizures: Expert consensus of the Epilepsy Foundation of America Working Group. *Epilepsia*, 46(9):1423–1425, 2005.
[29] J. Heer and G. Robertson. Animated transitions in statistical data graphics. *IEEE transactions on visualization and computer graphics*, 13(6):1240–1247, 2007.
[30] E. Hoque and M. Agrawala. Searching the visual style and structure of D3 visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):1236–1245, 2019.
[31] Y. Jansen, P. Dragicevic, P. Isenberg, J. Alexander, A. Karnik, J. Kildal, S. Subramanian, and K. Hornbæk. Opportunities and challenges for data physicalization. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pp. 3227–3236, 2015.
[32] C. Jung, S. Mehta, A. Kulkarni, Y. Zhao, and Y.-S. Kim. Communicating visualizations without visuals: Investigation of visualization alternative text for people with visual impairments. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):1095–1105, 2021.
[33] C. Kang. A tweet to Kurt Eichenwald, a strobe, and a seizure. now, an arrest. *The New York Times*, 2017.
[34] A. T. Le, J. Payne, C. Clarke, M. A. Kelly, F. Prudenziati, E. Armsby, O. Penacchio, and A. J. Wilkins. Discomfort from urban scenes: Metabolic consequences. *Landscape and Urban Planning*, 160:61–68, 2017.
[35] Y.-X. Lin, R. Venkatakrishnan, R. Venkatakrishnan, E. Ebrahimi, W.-C. Lin, and S. V. Babu. How the presence and size of static peripheral blur affects cybersickness in virtual reality. *ACM Transactions on Applied Perception (TAP)*, 17(4):1–18, 2020.
[36] Z. Liu and J. Heer. The effects of interactive latency on exploratory visual analysis. *IEEE transactions on visualization and computer graphics*, 20(12):2122–2131, 2014.
[37] A. Lundgard, C. Lee, and A. Satyanarayan. Sociotechnical considerations for accessible visualization design. In *2019 IEEE Visualization Conference (VIS)*, pp. 16–20. IEEE, 2019.
[38] A. Lundgard and A. Satyanarayan. Accessible visualization via natural language descriptions: A four-level model of semantic content. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):1073–1083, 2021.
[39] G. M. Machado, M. M. Oliveira, and L. A. Fernandes. A physiologically-based model for simulation of color vision deficiency. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1291–1298, 2009.
[40] K. Marriott, B. Lee, M. Butler, E. Cutrell, K. Ellis, C. Goncu, M. Hearst, K. McCoy, and D. A. Szafir. Inclusive data visualization for people with disabilities: a call to action. *Interactions*, 28(3):47–51, 2021.
[41] A. Matsumoto, S. Abe, T. Hiraki, S. Fukushima, and T. Naemura. Imperceptible AR markers for near-screen mobile interaction. *IEEE Access*, 7:79927–79933, 2019.
[42] T. Nightingale. Non-contiguous cartogram: Canada, Dec. 2015. http://bl.ocks.org/tnightingale/4704168/.
[43] A. Pearce. step-path-transition, Aug. 2016. http://bl.ocks.org/1wheel/7743519/.
[44] O. Penacchio, S. M. Haigh, X. Ross, R. Ferguson, and A. J. Wilkins. Visual discomfort and variations in chromaticity in art and nature. *Frontiers in neuroscience*, 15, 2021.
[45] O. Penacchio and A. J. Wilkins. Visual discomfort and the spatial distribution of fourier energy. *Vision research*, 108:1–7, 2015.
[46] K. Poulsen. Hackers assault epilepsy patients via computer. *Wired*, Mar. 2008.
[47] B. Rankin and M. Daniels. The shape of slavery. *The Pudding*, Jan. 2011.
[48] A. Satyanarayan, D. Moritz, K. Wongsuphasawat, and J. Heer. Vega-lite: A grammar of interactive graphics. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):341–350, 2016.
[49] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *The craft of information visualization*, pp. 364–371. Elsevier, 2003.
[50] A. Sielen. Reusable responsive multiline chart, June 2020. http://bl.ocks.org/asielen/44ffca2877d0132572cb/.
[51] G. Simoes. D3.js: Animating between scales, Apr. 2020. https://bl.ocks.org/guilhermesimoes/15ed216d14175d8165e6/.
[52] L. South and M. A. Borkin. Generating seizure-inducing sequences with interactive visualizations. 2020.
[53] L. South, D. Saffo, and M. A. Borkin. Detecting and defending against seizure-inducing gifs in social media. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pp. 1–17, 2021.
[54] H. Stevens. Basic map functions, Feb. 2018. http://bl.ocks.org/HarryStevens/1c07d73efaf074de05e63a33431eb80a/.
[55] J. Stucki. Swiss topography. *bl.ocks.org*, Dec. 2018. http://bl.ocks.org/herrstucki/6312708/.
[56] C. R. Systems. Harding flash and pattern analyzer.
[57] E. Vullings. Grouped horizontal bar chart., Aug. 2015. http://bl.ocks.org/erikvullings/51cc5332439939f1f292.
[58] M. Waldner, M. Le Muzic, M. Bernhard, W. Purgathofer, and I. Viola.

Attractive flicker—guiding attention in dynamic narrative visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2456–2465, 2014.

[59] P. Watson and A. Petrie. Method agreement analysis: a review of correct methodology. *Theriogenology*, 73(9):1167–1179, 2010.

[60] A. Wilkins, J. Emmett, and G. Harding. Characterizing the patterned images that precipitate seizures and optimizing guidelines to prevent them. *Epilepsia*, 46(8):1212–1218, 2005.

[61] J. Woodring and H.-W. Shen. Incorporating highlighting animations into static visualizations. In *Visualization and Data Analysis 2007*, vol. 6495, pp. 20–31. SPIE, 2007.

[62] K. Wu, E. Petersen, T. Ahmad, D. Burlinson, S. Tanis, and D. A. Szafir. Understanding data accessibility for people with intellectual and developmental disabilities. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pp. 1–16, 2021.

[63] K. Yourish, W. Andrews, L. Buchanan, and A. McLean. State gun laws enacted in the year after newtown. *The New York Times*, Dec. 2013.