

PixelClipper: Supporting Public Engagement and Conversation About Visualizations

Jagoda Walny

University of Calgary

Sarah Storteboom

University of Calgary

Richard Pusch

University of Calgary

Steven Munsu Hwang

University of Calgary

Søren Knudsen

University of Calgary and University of
Copenhagen

Sheelagh Carpendale

University of Calgary and Simon Fraser University

Wesley Willett

University of Calgary

Abstract—In this article, we present PixelClipper, a tool built for facilitating data engagement events. PixelClipper supports conversations around visualizations in public settings through annotation and commenting capabilities. It is recognized that understanding data is important for an informed society. However, even when visualizations are available on the web, open data is not yet reaching all audiences. Public facilitated events centered around data visualizations may help bridge this gap. PixelClipper is designed to promote discussion and engagement with visualizations in public settings. It allows viewers to quickly and expressively extract visual clippings from visualizations and add comments to them. Ambient and facilitator displays attract attention by showing clippings. They function as entry points to the full visualizations while supporting deeper conversations about the visualizations and data. We describe the design goals of PixelClipper, share our experiences from deploying it, and discuss its future potential in supporting data visualization engagement events.

Digital Object Identifier 10.1109/MCG.2020.2968906

Date of publication 23 January 2020; date of current version

28 February 2020.

■ **THE DESIRE TO** engage citizens in conversations grounded in data is a key factor in the open data movement in many democratic nations. Providing open data is seen as a path to transparent governance, and access to such data can increase citizens' power to understand and discuss important issues. This is not as simple as making data tables available online; even when data is openly available, citizens have disparate resources and skills to understand and engage with it. Many people do not have the time, tools, or skills to fully understand the data and its application to the decisions they make. To truly be accessible to the vast majority of citizens, governments must present the data in ways that are human-readable, engaging, and understandable.

An often-used strategy to make data more accessible is to use data visualizations. Visualizations exist to make data more understandable to people. They also tend to have broader appeal than raw data tables due to their graphical nature. However, even when visualized, data can be complex and difficult to interpret. As such, it is common for data visualizations to be presented with guidance and editorial filters. For example, Hans Rosling's guided presentations of data visualizations highlighting open global development indicators¹ are hugely popular, with over 14 million views on the TED website as of this writing. Journalists often use data-driven narrative techniques,^{2,3} such as explanatory text, visualization annotations, or stepped presentations of data visualizations to guide readers through datasets. These techniques are commonly used by graphics departments in news organizations such as the *New York Times* and *The Guardian*.

However, such guidance, by its nature, requires the author to make a judgement about which data points to highlight, what interpretations to suggest, and in some cases which conclusions to emphasize about the data. This presents a challenge in the case of government data, because in that context, presenting any particular point of view takes away from citizens' ability to use the data to form their own conclusions. While it is impossible to be completely neutral when collecting or presenting data, it is important for governments and

other ostensibly unbiased institutions to minimize the imposition of interpretations on the audience. This poses a challenge for engaging citizens with open data, because there is a need to balance neutral presentation of data with provision of the necessary support for people to be able to interpret it.

Holding *in-person events* around data visualizations is one way to address this challenge. In-person events are events held locally in publicly accessible physical spaces that are facilitated by data experts present at the event. Social interaction around data visualization has been identified as one way to promote better understanding.⁴ While this has largely been investigated from the standpoint of distributed online collaborative analysis of data via tools such as ManyEyes,⁵ Sense.us,⁶ and CommentSpace,⁷ our focus is on in-person events. While smaller in scale than online tools, in-person events can provide deep social interaction around data. Such events provide an opportunity for individuals to become curious about the data and to ask for specific one-on-one guidance that aligns with their interests and needs. While such events reach a far smaller audience than an online data visualization can, they provide an opportunity to reach people who might not otherwise engage with a web visualization as deeply, or at all. Ideally, a public data visualization engagement event can not only raise awareness of the open data in question, but also provide opportunities for members of the public to talk to data experts—the people who work with the data on a regular basis and who can provide answers to questions and tailored guidance about how to navigate and learn from the data visualizations. An event can also provide a supportive, reflective environment in which engagement with data can occur.

We report on a digital system called Pixel-Clipper, which we developed to support in-person facilitated public data visualization engagement events. In our experience, presenting data visualizations in a public context without tool support tends to encourage a facilitator-led model. Usually, a facilitator guides visitors through set views in the visualization, but only a minority of visitors interact with the visualizations themselves. The main

function of PixelClipper is to provide a platform to encourage event visitors and facilitators to interact with the visualizations and hold conversations about specific data visualization views. PixelClipper allows visitors and facilitators alike to create comments that serve as entry points into the event and into in-person conversations. We designed PixelClipper to integrate with our own engagement events where we showcase a set of open data visualizations we designed during a multiyear project with the Canada Energy Regulator (CER).*

The PixelClipper system (see Figure 2) includes a set of tools that allow visitors to extract, share, and annotate clippings of web-based visualizations during public visualization events. These clippings act as entry points into various different points in the visualization and both viewers and facilitators can use them as jumping-off points to facilitate richer conversations and exploration. The system consists of an ambient display for showcasing clippings, a facilitated display for in-person discussion of clippings and visualizations, a clipping tool, and a website for storing and sharing clippings online.

OPEN ENERGY DATA VISUALIZATIONS

We created PixelClipper within the context of a multiyear project during which we designed a set of data visualizations of open data in collaboration with the CER, a Canadian government organization. The goal of the project was to make the data not only available, but also accessible to Canadians. As part of this objective, we also organized and coordinated several engagement events facilitated by data experts. The events featured four visualizations of open data from the CER (see Figure 1). The first, *Energy Futures*, visualizes projections about energy production, use, and demand in Canada until the year 2040. The second, *Pipeline Incidents*, visualizes reported incidents on CER-regulated pipelines and pipeline facilities across Canada. A third, *Imports and Exports*, visualizes the imports and exports of energy projects into and out of Canada. Finally, *Pipeline Conditions* visualizes the regulatory conditions placed on CER-

regulated pipeline projects. All visualizations are publicly hosted on the website of the CER.[†] The CER provides and maintains all of the data used in these visualizations and the data experts who facilitated the events we describe are employees of the regulator and actively work with the visualized datasets.

DESIGNING FOR FACILITATED DISCUSSIONS

Complex visualizations designed for the web can be challenging to demonstrate in public settings. In our experience across several decades of visualization demos, presenters are usually predisposed to present visualizations' features and technical capabilities, rather than providing opportunities for viewers to engage deeply with the data. Furthermore, this facilitator-led model gives implicit control over interaction to the facilitator, so that visitors rarely interact with the data visualizations directly. Viewers rarely have opportunities to interact with the visualizations for extended periods, and the public nature of the interaction can create social dynamics that discourage them from exploring or engaging deeply.

With this premise, we examined opportunities for engagement events that allowed visitors to play a more active role. We considered initial alternatives: 1) displaying visualizations in a central place at the event as we had done in previous events, and 2) letting visitors explore the visualizations and content on their own using workstations. We then considered how to connect these different aspects of the event in a way that could support transitions between different levels of engagement with the visualizations.

A prevailing definition of engagement with visualizations does not yet exist, though it is beginning to be explored in the visualization research community. Hung and Parsons⁸ suggest that research into visualization engagement should consider factors such as intention, autonomy, purpose, time, and outcome. Meanwhile, Mahyar *et al.*⁹ use cognitive task level as a proxy for engagement, where viewing and interacting

*Formerly the National Energy Board of Canada.

[†]<http://www.cer-rec.gc.ca>. Redirect links are also available at <http://ilab.cpsc.ucalgary.ca/energyvis>

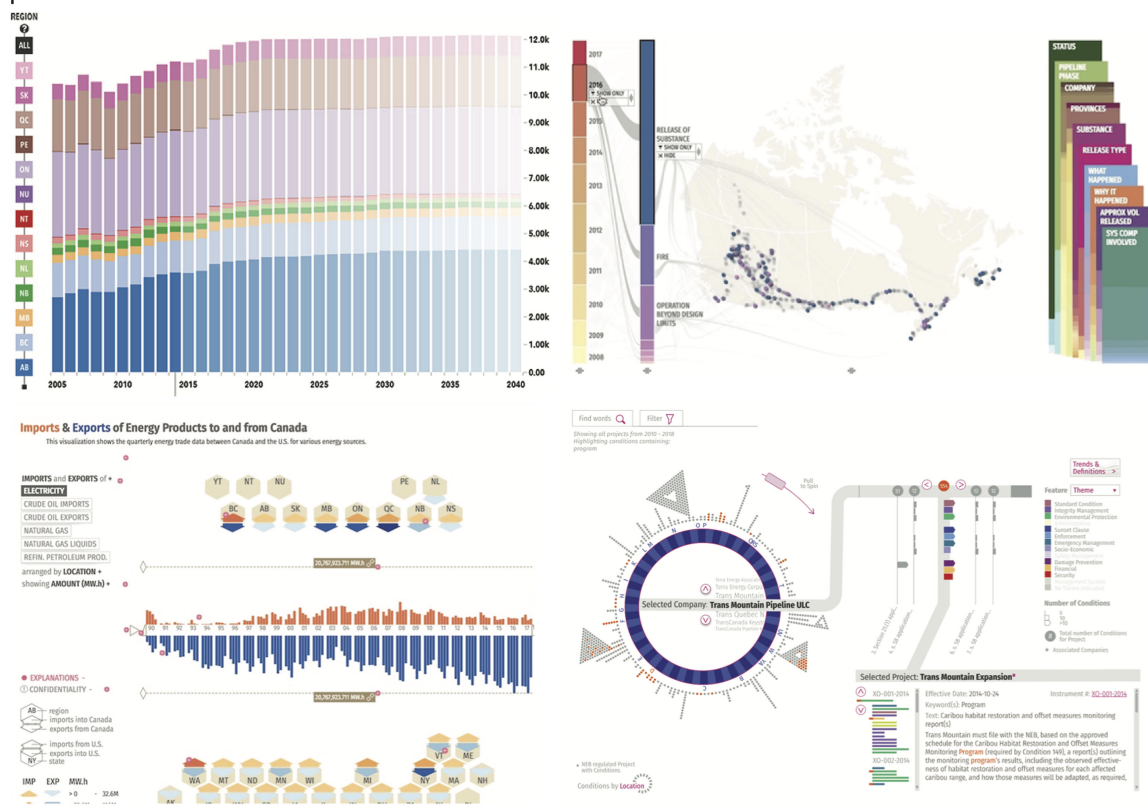


Figure 1. Views from each of the four data visualizations around which we based our data visualization engagement event design. Clockwise from top left: Energy Futures, Pipeline Incidents, Pipeline Conditions, and Energy Imports/Exports.

are considered as low engagement while synthesizing and deriving decisions are considered as high engagement. In contrast, we focus primarily

on two specific manifestations of engagement in the context of public events—*direct engagement* with the visualization and facilitators, and *engagement with data* via exploration and conversation (see Table 1). These dimensions provide a lens with which to characterize participants' interactions with the events themselves and their connection to the data that the visualizations showcase.

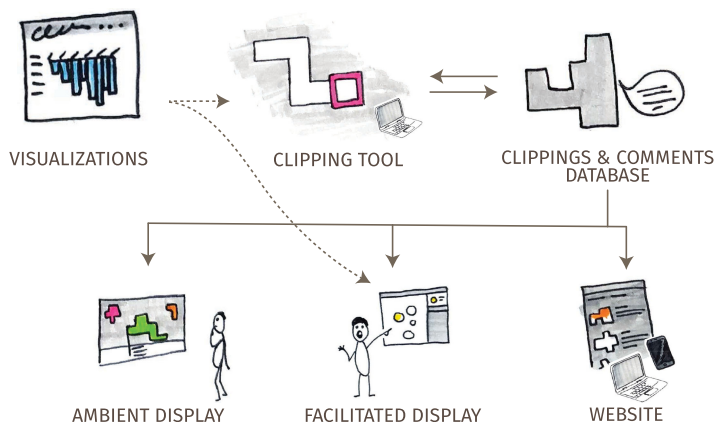


Figure 2. The complete PixelClipper system includes (a) the clipping and commenting tool, (b) a backend server with the clipping and comments database and administrator interface, (c) an ambient display interface, (d) a facilitator display interface, and (e) a mobile-friendly website for displaying clippings. Solid arrows show the flow of clipping data within PixelClipper, whereas dashed lines indicate where the visualizations can be accessed in full.

Direct Engagement is characterized by a visitor's observable interactions with the space, artifacts, and people at the event. We place direct engagement actions along a spectrum that includes passive engagement on the low end, active engagement in the middle, and return on the high end. *Passive* engagement occurs when a visitor observes but does not affect the artifacts or people at the event—for example by glancing at a display or listening to a facilitator. *Active* engagement occurs when a visitor affects the artifacts or people at the event—for example by interacting with a digital system or conversing with a

facilitator. *Return* engagement occurs when a visitor returns to the event itself or any aspect of the event via dialogue or via physical and digital artifacts after the event has taken place.

Data Engagement is characterized by the extent to which a visitor connects with the core content of the event—in this case, the data behind the visualizations. This is largely an internal process and therefore not always observable or directly measurable. However, it remains a key goal of a data visualization engagement event. We characterize low data engagement as simple *awareness* of the data's existence. Medium data engagement occurs when a visitor actively *explores* the data and begins to understand its structure and meaning. High data engagement occurs when a visitor *reflects on or responds* to the data. For example, a highly engaged visitor might come to a conclusion about the data, form an opinion of it, decide how to apply it to his or her life, pose questions about the data, or comment about it to a facilitator or other visitors.

Design Goals

With these dimensions in mind, we set out to create a tool that could support public discourse and engagement with open data visualizations at public events. In particular, we sought to create a tool that would allow diverse audiences to engage with data visualizations both privately and via public conversations, while creating opportunities for more direct interaction with facilitators and other visitors as well as deeper engagement with the data. To this end, we articulated the following specific design goals:

DG1: Support Multiple Degrees of Direct Engagement With the Visualizations

Direct engagement speaks to the overall attraction and appeal of an event. The amount of direct engagement at the event can affect the perceived atmosphere and energy of the event, which can have an impact on attracting more visitors and encouraging data engagement. Furthermore, while direct engagement does not guarantee data engagement, some form of direct engagement is a prerequisite to engaging with the data. Moreover, direct interaction with the

Table 1. Dimensions of engagement.

Engagement level	Low	Medium	High
Direct engagement	Passive	Active	Return
Data engagement	Awareness	Exploration	Response

tools and people present can itself be an entry point into the data. The tool should support passive engagement to attract visitors and serve as an entry point into the event, encourage active engagement during the event, and facilitate a return to the content at later date for those who are so inclined.

DG2: Support Multiple Degrees of Engagement With the Data

The ultimate intent of a data visualization engagement event is to engage visitors with the underlying data. Different visitors will be inclined toward different levels of engagement. Some will want to deeply explore the data while others may be content to simply be aware that it exists. The event succeeds in its intent if it engages visitors at any point on the data engagement spectrum, not just the high end. The tool should be sufficiently flexible to support those who have little knowledge or awareness of the data and those who are inclined to explore the data more deeply. It should also support the highest degree of data engagement, including reflection on and response to the data, whether this be at an individual level or in conversation.

DG3: Support Both Independent and Facilitated Exploration of the Visualizations

The ability to bring expert facilitators and visitors together in one place is a key advantage of the in-person event context. However, supporting individual exploration can be just as important. Some visitors are often more inclined to engage with the data independently, while others may want to switch between individual and facilitated exploration. As such, the tool's design should support facilitators in guiding visitors through the data visualizations while simultaneously supporting those interested in engaging with visualizations independently.

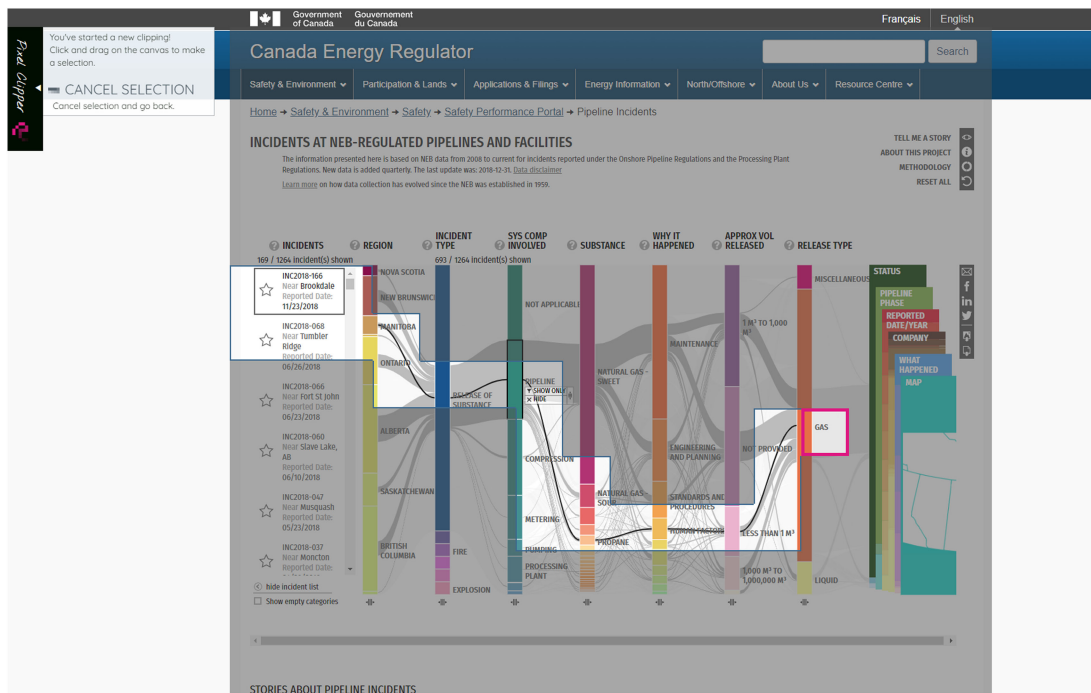


Figure 3. PixelClipper's clipping tool interface in clipping mode shown on top of a filtered view of a web visualization.¹⁰ The white shape consisting of contiguous grid blocks is being selected using a single mouse drag gesture. The pink square indicates the current location of the mouse cursor.

PIXELCLIPPER

With these goals in mind, we developed PixelClipper, a system that allows visitors to clip and annotate visualization views and then use them to guide conversations with data experts. The system (see Figure 2) consists of an expressive clipping and commenting tool, a facilitation interface for reading comments, an ambient display to give visitors a preview of the comments, and a mobile-friendly website for browsing and sharing clippings on social media.

We designed PixelClipper to work with a variety of existing web-based data visualizations, including those we developed for the CER. The system relies on visitor- and facilitator-created clippings to provide *entry points* into the visualization that support multiple types of engagement.

Clipping and Commenting Tool

The clipping and commenting tool (see Figure 3) is inspired by the analog act of cutting out clippings from magazines or newspapers to save specific pieces of information. It uses a desktop browser extension that allows both visitors and facilitators to isolate and annotate portions of a visualization view. Clippings can be

accompanied by short text comments and link back to the view they were clipped from. As such, the clipping tool provides: (a) an alternative approach to individually exploring the original web-based data visualizations; (b) a platform for sharing reflections, questions, and insights about the data; and (c) a way for facilitators to start a conversation with attendees. The clipping tool is the backbone of PixelClipper and the event, as it provides a way for facilitators and visitors to create the clippings and comments that will be displayed on all other parts of the system.

When a visualization is viewed with the extension activated, the extension displays a sidebar of commands that include controls for creating a clipping and browsing the complete database of clippings made by others. The tool also supports “doubly-linked discussion”⁶ via a collapsible panel that shows other clippings made on the same visualization view—making it easier for visitors to see others’ contributions as they browse.

The mechanics of creating a clipping reflect a careful balance of expressiveness, precision, and simplicity. The extension overlays a coarse



Figure 4. Viewers create clippings by drawing a path on top of the visualization. This highlights a set of cells in the selection grid for inclusion in the clipping and also determines their ordering in future animations.

two-dimensional grid on top of the visualization. With one fluid mouse drag, the visitor can select a portion of the visualization to “clip” (see Figure 4). Once they release the mouse cursor, a text box appears, allowing the visitor to associate a brief comment with the clipping. This deliberately simple selection gesture limits the number of overall clicks and choices that are required to make a selection, keeping most of the focus on the visualization. The underlying coarse grid system provides some expressive freedom in how the shapes are drawn, but not so much as to make the clipping tool feel like a drawing program. The commenting interface supports brief text comments, which can be either anonymous or can include a commenter’s name or pseudonym.

Before posting clippings or comments, the tool asks visitors to give consent, allowing that content to be stored, displayed, and (optionally) used in research. A database on a central server stores each clipping image, comment text, and attribution. It also stores the unique URL of the original visualization view, making it possible to return to that exact visualization state in the future. These clippings are then displayed in the facilitation interface and the ambient display, which are described in the following sections.

Implementation. We implemented the clipping tool as a Google Chrome extension and opted to make the extension publicly downloadable in addition to providing fixed computer stations during events. While the tool can technically be used to clip or comment on any webpage, we restrict the extension to work only on a set of pre-approved web-based visualizations that provide unique, stateful URLs for individual visualization views and which we have cleared for participation in the program.

Engagement. The clipping tool is the key piece of PixelClipper that gives visitors a chance to engage with the event in multiple ways. Active engagement happens through the act of exploring the visualization and creating a clipping. At the same time, the clipping tool generates content that can serve as entry points for other visitors. It also provides content that visitors can return to in future interactions with the visualizations and with other pieces of the PixelClipper system.

The clipping tool is intended to support the deepest levels of engagement with respect to all three design goals. It maintains the ability for full interaction with the basic content of the event, that is, the data visualizations. It provides an opportunity for visitors to learn to navigate the visualizations and potentially return at a later date (DG1). Its expressive interface can be rewarding in and of itself, regardless of the visitor’s engagement with the data, inherently encouraging direct interaction (DG1). It also provides a simple way to express reflections, thoughts, and questions about the data (DG2). Finally, it allows facilitators to point out areas of interest within the visualization and seed later discussions by adding their own, specially marked comments (DG3).

Facilitation Interface

One of the biggest advantages of in-person events is that facilitators can guide visitors through a visualization and answer their questions (DG3). PixelClipper’s facilitation interface (see Figure 5) allows facilitators and visitors alike to view existing clippings and comments and return directly to the same data visualization views from which the clippings were made. We designed the facilitation interface primarily to support facilitators in conversations with event attendees around large touch-enabled displays.

The left side of the facilitation interface, shown in Figure 5, shows fully-interactive versions of the visualizations themselves. A sidebar on the right displays a list of clippings for the selected visualization, ordered from newest to oldest. Each shows the clipping image in its entirety together with a comment, date, and author. Facilitator comments are marked

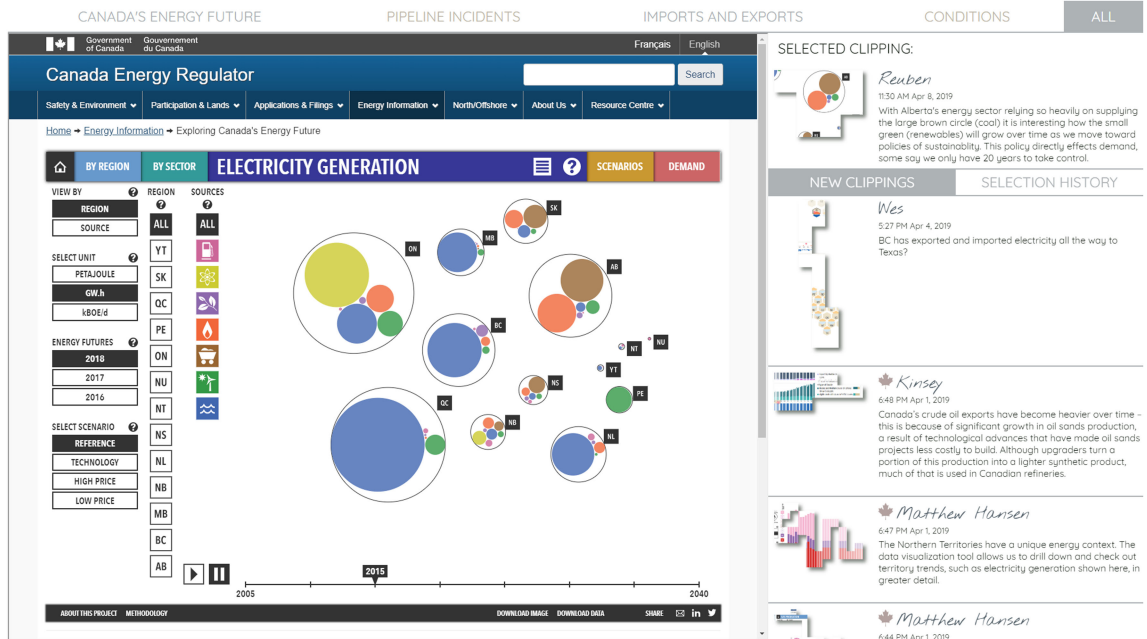


Figure 5. PixelClipper's facilitation interface showing clippings from the Energy Futures visualization.¹⁰

with an icon (a maple leaf, which references that our data experts are Canadian federal government employees). When a clipping is selected, the visualization frame loads the view of the visualization that is associated with that clipping. Meanwhile, that clipping is moved to the “selected clipping” pane in the top right corner and is added to the “selection history” view in the sidebar. The facilitation interface does not itself support the creation of new clippings.

Implementation. The facilitation interface reads clippings from the clippings database on a central server and loads the visualizations into an iframe. It is primarily designed for 1080p large touch displays.

Engagement. PixelClipper's facilitation interface stems directly from DG3, supporting both individual and facilitated exploration. It is meant to cater mainly to people with medium to high direct engagement who are interested in what the facilitators and other visitors are saying about the data. Facilitators can use this interface to drive their conversations with visitors. For those visitors who are less engaged with content, facilitators can “seed” the clippings database with specific predetermined entry points that might spark interest or provide some clarity about the contents of the visualization. At the

same time, they can “read the room” by seeing all of the comments and questions that visitors are adding to the database, thereby tailoring their conversations to the level of engagement that visitors have.

Ambient Display

The final piece of PixelClipper, the ambient display (see Figure 6), is designed to attract potential visitors to the event space through passive interaction and awareness. The ambient display continuously shows images of the most recent visualization clippings. These clippings slowly animate in and out atop a dark pixelated background that evokes the selection grid used by the clipping tool. As they appear, each clipping animates in using the same order in which the author marked the blocks in the clipping tool. This emphasizes the clipping author's process and provides a subtle avenue for further personalization of the clippings. Clippings first appear as solid colored blocks, then reveal the clipped slice of the visualization, along with any associated comments. After a short pause, the clippings then fade back to a colored block and disappear. The initial solid color creates a sense of mystery and anticipation for the reveal of the selection.

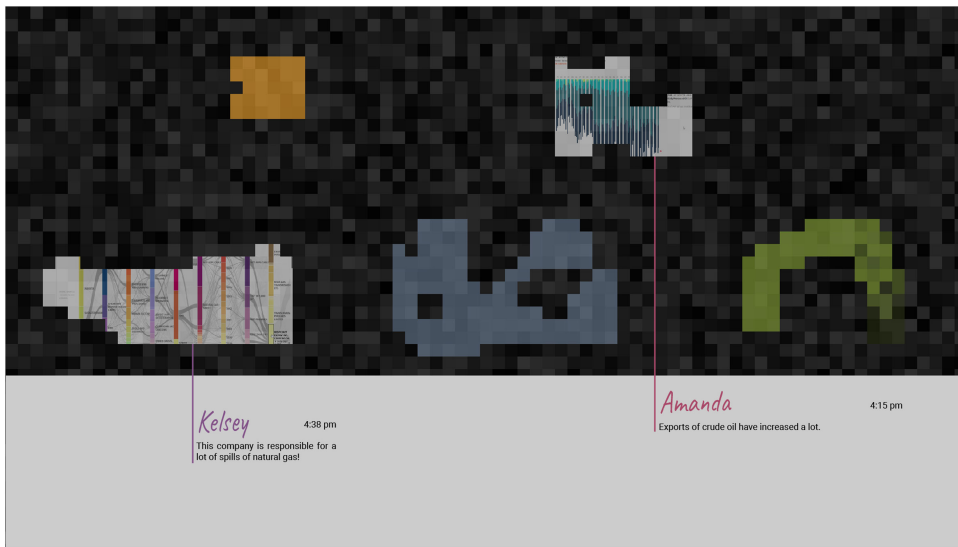


Figure 6. View of PixelClipper's ambient display showing a selection of clippings for all of the visualizations used in the event, in various stages of animation.

Implementation. Like the facilitator display, we implemented the ambient display in JavaScript and designed it to run in a web browser. The program relies primarily on animation of the unique clipping shapes to attract viewers. The ambient display shows a semi-random subset of clippings through a slow and continuous animation. However, to balance visitor- and facilitator-created comments, the display injects facilitator comments at regular intervals.

Engagement. The ambient display program partially satisfies DG1 and DG2 by providing an accessible entry point for visitors who are unaware of the visualizations or the data and may need additional incentive to engage with the event. Slowly revealing a subset of clippings, rather than showing full visualization views, is a deliberate attempt to pique visitors' interest, stimulate reflection, and attract their attention to the event as a whole.

Mobile-Friendly Website

PixelClipper also includes a mobile-friendly website that can display all of the clippings in the database on a visitor's personal device both during and after the event. It also contains links to the original data visualizations and supports sharing clippings on social media. This satisfies the *return* end of the direct engagement dimension.

Support Program

Finally, PixelClipper incorporates an administration interface that supports basic administrative tasks such as moderation of comments and viewing of the full database by event organizers.

EVENT CONTEXT

Alongside PixelClipper we have designed a facilitated data visualization engagement event plan that uses all of the components (see Figure 7) of the PixelClipper system. The various pieces of the system can work flexibly in the context of different event spaces and with available technology. Our experience and observations deploying an early version of PixelClipper at a pilot event had a large influence on the current design.

Pilot Event

In April 2019, we held a pilot event (see Figure 8) in the lobby of a university library using an early version of PixelClipper. The event lasted a full day and showcased three of the visualizations—Energy Futures, Pipeline Incidents, and Imports and Exports (the Pipeline Conditions visualization was not yet available). The event was advertised through the University and CER networks, but primarily was geared toward spontaneous attendance by

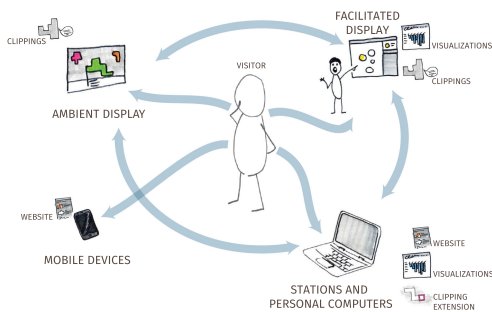


Figure 7. Visitors can flow between all components of the PixelClipper system within the event space and take away a mobile experience with the website.

visitors drawn from the roughly five thousand passersby who transit the library's lobby each weekday. Two data experts were present as facilitators and several members of the visualization design team were also present to provide interaction guidance.

The event used two large touchscreen displays that, in this prototype of PixelClipper, doubled as both ambient and facilitation displays. On these screens, the facilitation interface could be accessed directly from the ambient display. Selecting a clipping in the ambient display launched a window in the middle of the display showing the corresponding visualization view. Facilitators were limited in that they could only open visualizations from the subset of clippings currently being displayed. When the visualization window was open, the selected clipping and comment were often obscured.

Our event space included seated workstations for individual exploration and use of the clipping tool, which was not yet available for download on private devices. We recorded video from the event for research purposes. To do so, we arranged the displays to create a contained recording zone that also included the seated stations. The wall behind the event space featured a large display that we used as a static banner advertising the event.

Observations

After the event, we held discussions with the members of the research team about their experiences and observations during the event. We found that the contained recording zone created a clear inside/outside space of the event and that some individuals who appeared mildly



Figure 8. PixelClipper deployed as part of a public data outreach event. Top: Displays in a library thoroughfare. Left: Discussion using an early hybrid ambient display/facilitation interface. Right: Clipping and commenting at a capture station.

interested seemed uncomfortable crossing the threshold. However, when visitors did enter the space, both the facilitated displays and the individual exploration stations were well-used. Some visitors spent over half an hour at the seated stations and created multiple clippings that indicated a deep level of data engagement.

Both the ambient display and facilitation views were valuable during the event. However, because the two were combined on the same displays, they often competed for attention. The design of the interface also became over-complicated because, while the ambient display ideally would be simple and eye-catching, we also wanted to provide the facilitators with options to guide their discussion. These two functions were in conflict with one another when sharing a display. This was the main motivation for separating the facilitation and ambient displays in the current iteration.

Finally, we found that our initial event failed to support return engagement for individuals who were not able to stay at the event but were eager to engage with the content. We had postcards printed with the main website address for the visualizations, but because the visualizations are not compatible with mobile devices, individuals would have to wait until they were at a desktop computer or tablet before returning to the visualizations. This was the impetus for creating the

mobile-friendly website in the most recent version of PixelClipper, which allows visitors to engage with the clippings on mobile devices, even if the visualizations themselves cannot be viewed.

Current Setup

In our more recent event configurations, we position the components of the PixelClipper system that facilitate low direct interaction at the foreground and those that support high direct interaction at the center. We carefully position the ambient displays in visible locations at the front of the event space, targeting visitors who are passively engaged. These screens aim to capture attention from high traffic areas and areas where individuals tend to linger when resting or waiting. We also use the ambient displays to direct people to our website, where they can browse the clippings on any device. Meanwhile, we position the facilitator displays in a central position where they can support both passive and active engagement. Finally, we include multiple computers in the space in both standing and seated configurations. These are loaded with the clipping tool. We position the standing stations between the ambient displays and the facilitator displays to allow for low-commitment interactions. In contrast, we position the seated stations in the background for longer-term browsing and exploration by visitors with high direct and content engagement. We have also made it possible for event organizers to make the PixelClipper extension publicly available to install on personal devices, creating the opportunity for it to be used on personal computers before, during, and after an event.

DISCUSSION

Throughout our prototyping and initial deployment of PixelClipper, we have engaged in an ongoing process of reflection. This provided an opportunity to contrast the experience of using PixelClipper against our other experiences of exhibiting visualizations in public settings and highlighted a number of challenges and opportunities for future facilitation tools.

In particular, our initial deployments have emphasized the extent to which PixelClipper supports longer sustained interaction with visualizations than is typical in public settings. In

comparison to other demos we have run using the same visualizations, visitors at our deployment at the library spent much longer engaging with the content, with several passersby ultimately spending more than 30 min exploring the visualizations and creating clippings. We suspect that the interactions with facilitators and the relative privacy of the browsing stations gave participants explicit permission to engage with the visualizations, unlike typical software demonstrations in which the facilitator has implied ownership over the interaction space. Furthermore, the simple, constrained interface of PixelClipper gave participants an opportunity to contribute even before they had learned to operate the visualization. Meanwhile, the event still succeeded in attracting other kinds of engagement, including individual interactions with the ambient display and with facilitators. However, our experiences designing and deploying PixelClipper also suggest potential challenges and opportunities for facilitated discussion tools.

Decontextualization of the Visualizations

through the way in which PixelClipper's clippings are extracted from the original visualizations presents tradeoffs. The main drawback is that most freeform clippings lose some of the context of the visualization, and many omit elements like titles, legends, filters, and transition animations that are often essential for interpreting a visualization. However, these elements, while important, create additional cognitive load for viewers. The decontextualized design of the clippings allows authors to create simpler snippets that distill visually interesting or important elements of the visualization without overwhelming viewers. The primary function of the clippings, supporting design goals DG1 and DG2, is to attract visitors to the content of the visualizations even if their initial level of engagement is quite low. Much like headlines in a newspaper, the clippings only hint at the full scope of the data, and act as points of entry into the visualization or related discussions for those whose interest is piqued. However, as with newspaper headlines, there remains a risk that the lack of context can lead to misinterpretations. It is also possible that some visitors may be turned away

by the lack of interpretive guidance, particularly when authors do not annotate their clippings. In a public event context, some of these risks can be mitigated by the presence of facilitators and by including multiple displays showing different views of the data.

Obscure Entry Points created when viewers author new clippings present an opportunity for others to enter a dataset at personally relevant locations, encouraging engagement with the data. Comments and clippings can highlight interesting and hidden points in a visualization and provide entry points that can help make a visualization or dataset feel relevant to a wider audience. However, creating more entry points also gives visualization designers less control over viewers' first experience of the visualization and can make it harder to provide context or introduce important pieces of functionality. In PixelClipper, we balanced this to a degree by ensuring that facilitator-created clippings appear regularly on the ambient display. However, providing customized introductions to visualizations based on entry points is an interesting opportunity for future work.

Facilitator Liability can also pose a challenge for open-data discussions, since it can limit the kinds of responses facilitators are able to give. For complex datasets like the ones we considered, government employees and data experts are uniquely positioned to explain and respond to questions from the public. However, individual employees may not have the permission to give official statements on behalf of their larger organization. This can limit the degree to which these employees can serve as facilitators and the kinds of responses they can give. In our case, employees from the CER were not permitted to respond to viewers' questions in writing within the commenting system, which limited their ability to provide persistent answers or clarifications. Supporting fully-fledged discussions with these kinds of data experts requires buy-in from organizations who can allocate resources to crafting these responses and obtaining legal clearance to release public statements on sensitive topics.

Vandalism and Abuse are possible in any public forum, and are particularly likely to arise in discussions around politically-charged topics. Thus far, we have seen little evidence of abuse in our deployments and have reviewed and moderated comments on an as-needed basis. Proactive moderation may be necessary in more heated contexts, but would introduce delay and moderation costs. Moreover, it seems clear that larger public deployments will need to deal with the same issues of scale and judgement that plague content moderation in online settings.

Bridging Events With Online Discussion represents an exciting future opportunity for future facilitated visualization tools. While systems like PixelClipper highlight the value of in-person discussion around visualizations, they also suggest how those discussions could connect to the kinds of broader online discussion that occur on news sites and social media. For example, clipping and commenting tools like ours could easily connect to existing social media platforms like Instagram or Twitter, providing new ways of annotating and publicly sharing these kinds of visualizations. While we have incorporated this into PixelClipper's website component, integrating social media sharing into the main public-facing event pieces could impact the experience of the event and potentially open up participation in the event to nonlocal audiences. Similarly, ambient displays and facilitator tools could also be used to facilitate in-person public discussions that build off of ongoing discussions online.

Persistent Public Installations in settings like libraries and museums could also leverage these kinds of tools to support richer low-level data exploration and engagement. For example, ambient clipping and comment displays like those we created could serve as a nexus for longer term engagement and commenting around datasets at a single site. Venues could also pair these persistent displays with periodic facilitated discussions that build off of the ongoing digital discussion. However, creating these kinds of experiences may call for additional functionality like voting or curation tools that would allow

organizers and facilitators to prioritize comments and clippings for discussion.

Expressive Clipping and Annotation Tools like those introduced in PixelClipper offer a rich yet simple mechanism for sharing snippets of visual content and may have promise in other settings. This clipping metaphor allows commenters to highlight important subsets of a visualization using a single stroke, while simultaneously using the direction of the stroke to direct viewers' attention. Clippings created using the tool also have distinct, recognizable, and expressive shapes that may make them more visually salient than more traditional rectangular clippings. This clipping metaphor could also easily be applied to other kinds of media including images and video.

The Complexity of Factors That Contribute to Engagement leads to limitations in our ability to fully understand the impact of a system like PixelClipper. The system is designed for in-the-wild events where awareness of data collection can influence how individuals experience and react to the event. Many aspects of engagement are not directly observable, but collecting feedback on the experience of transient event visitors presents major practical and ethical challenges. In addition, engagement with data and visualizations can be impacted by more factors than it is possible to address with a system like PixelClipper. These factors can include, for example, individuals' previous knowledge of and interest in content or facilitators' experience and approaches to facilitation. The lack of control over these factors in a public event makes it difficult to exactly measure the impact of PixelClipper. Our reflection on our experiences and design choices provides a starting point for such deeper investigations of engagement with data at public events.

CONCLUSION

While data visualizations are often demonstrated and presented in public spaces, these demonstrations still rarely result in deep engagement with data. Our PixelClipper prototype showcases the potential of expressive clipping and commenting tools to create new entry

points into visualizations and provide viewers with lightweight ways of contributing to data-driven discussions. Our system also illustrates how ambient and facilitator displays can leverage these contributions to create interest and support richer in-person conversations grounded in the data. Based on our initial experiences, these kinds of facilitation tools hold considerable promise and might support public engagement with a diverse range of datasets. Moreover, we believe that these tools may serve as building blocks for an even wider range of systems that bridge online and offline discussions around data.

ACKNOWLEDGMENTS

This work was supported in part by the Canada Energy Regulator, in part by the National Sciences and Engineering Research Council of Canada, in part by Alberta Innovates Technology Futures, in part by SMART Technologies, and in part by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant agreement No. 753816.

REFERENCES

1. H. Rosling, "Hans Rosling: The best stats you've ever seen," [Video file], Feb. 2006. [Online]. Available: https://www.ted.com/talks/hans_rosling_the_best_stats_you_ve_ever_seen
2. E. Segel and J. Heer, "Narrative visualization: Telling stories with data," *IEEE Trans. Vis. Comput. Graphics*, vol. 16, no. 6, pp. 1139–1148, Nov./Dec. 2010.
3. N. Henry-Riche, C. Hurter, N. Diakopoulos, and S. Carpendale, *Data-Driven Storytelling*. New York, NY, USA: CRC Press, 2018.
4. J. Heer, F. B. Viegas, and M. Wattenberg, "Voyagers and voyeurs: Supporting asynchronous collaborative information visualization," in *Proc. ACM CHI*, 2007, pp. 1029–1038.
5. F. B. Viegas, M. Wattenberg, F. van Ham, J. Kriss, and M. McKeon, "ManyEyes: A site for visualization at internet scale," *IEEE Trans. Vis. Comput. Graphics*, vol. 13, no. 6, pp. 1121–1128, Nov./Dec. 2007.
6. J. Heer and M. Agrawala, "Design considerations for collaborative visual analytics," *Inf. Visualization*, vol. 7, no. 1, pp. 49–62, 2008.

7. W. Willett, J. Heer, J. Hellerstein, and M. Agrawala, "Commentspace: Structured support for collaborative visual analysis," in *Proc. ACM CHI*, 2011, pp. 3131–3140.
8. Y.-H. Hung and P. Parsons, "Assessing user engagement in information visualization," in *Proc. ACM CHI Extended Abstr.*, 2017, pp. 1708–1717.
9. N. Mahyar, S.-H. Kim, and B. C. Kwon, "Towards a taxonomy for evaluating user engagement in information visualization," in *Proc. IEEE VIS Workshop Pers. Visualization: Exploring Everyday Life*, 2015, pp. 40–43.
10. Canada Energy Regulator, "CER—Data visualization," Accessed: Jan. 2020. [Online]. Available: <https://www.cer-rec.gc.ca/dv/index-eng.html>

Jagoda Walny currently leads the Visual Data Group with the Canada Energy Regulator, Calgary, AB, Canada. She has recently completed a postdoctoral fellowship at the University of Calgary studying the design of public-facing data visualizations. She received the Ph.D. degree from the University of Calgary in 2016. Contact her at jkwalny@ucalgary.ca.

Sarah Storteboom is currently a Designer with the University of Calgary, Calgary, AB, Canada. She received the M.Sc. degree in computational media design from the University of Calgary and a BFA degree from Emily Carr University of Art and Design, Vancouver, BC, Canada. Contact her at sarah.storteboom@ucalgary.ca.

Richard Pusch is currently a Developer and Designer with the University of Calgary, Calgary, AB, Canada. After receiving the M.Sc. degree in computer graphics, he focused on designing and implementing interactive systems that allow people to create, share, and tell visual stories with data. Contact him at rapusch@ucalgary.ca.

Steven Munsu Hwang is an undergraduate student in computer science with the University of Calgary, Calgary, AB, Canada. Contact him at munsu.hwang@ucalgary.ca.

Søren Knudsen is currently a Marie Curie Postdoctoral Research Fellow with the University of Calgary, Calgary, AB, Canada. He leads the HIVElab, a cross-faculty team in the intersection of information visualization, human–computer interaction, and data science, focused on seeking ways to engage healthcare professionals and patients in meaningful conversations around data. He received the MSE degree from the Technical University of Denmark, Lyngby, Denmark, in 2008, and the Ph.D. degree from the University of Copenhagen, Copenhagen, Denmark, in 2015. Contact him at sknudsen@ucalgary.ca.

Sheelagh Carpendale is currently a Full Professor with Simon Fraser University, Burnaby, BC, Canada, where she leads her research group InnoVis (Innovations in Visualization) and the ixLab (Interactive Experiences Lab). She was the recipient of the IEEE Visualization Career Award and has been inducted into the ACM CHI and IEEE VIS Academies. Her interdisciplinary approach to research focuses on information visualization, human–computer interaction, personal visualization, and visualization to support data engagement and empowerment. Contact her at sheelagh@sfu.ca.

Wesley Willett is currently a Canada Research Chair in Visual Analytics with the University of Calgary, Calgary, AB, Canada. He leads the Data Experience Laboratory, an interdisciplinary research group whose research spans information visualization, human–computer interaction, and data science and focuses on pairing data and interactivity to support collaboration, learning, and discovery. Contact him at wesley.willett@ucalgary.ca.