




Extending ggplot2 for Linked and Animated Web Graphics

Carson Sievert, Susan VanderPlas, Jun Cai, Kevin Ferris, Faizan Uddin Fahad Khan & Toby Dylan Hocking


To cite this article: Carson Sievert, Susan VanderPlas, Jun Cai, Kevin Ferris, Faizan Uddin Fahad Khan & Toby Dylan Hocking (2019) Extending ggplot2 for Linked and Animated Web Graphics, Journal of Computational and Graphical Statistics, 28:2, 299-308, DOI: [10.1080/10618600.2018.1513367](https://doi.org/10.1080/10618600.2018.1513367)

To link to this article: <https://doi.org/10.1080/10618600.2018.1513367>

 View supplementary material [↗](#)

 Published online: 14 Nov 2018.

 Submit your article to this journal [↗](#)

 Article views: 499

 View related articles [↗](#)

 View Crossmark data [↗](#)

 Citing articles: 1 View citing articles [↗](#)



Extending ggplot2 for Linked and Animated Web Graphics

Carson Sievert^a, Susan VanderPlas^a, Jun Cai^b, Kevin Ferris^c, Faizan Uddin Fahad Khan^d, and Toby Dylan Hocking^e

^aDepartment of Statistics, Iowa State University, Ames, IA; ^bDepartment of Earth System Science, Tsinghua University, Haidian Qu, Beijing Shi, China; ^cBaseball Operations Department, Tampa Bay Rays, St. Petersburg, FL; ^dDepartment of Computer Science & Engineering, IIT BHU, Varanasi, Uttar Pradesh, India; ^eDepartment of Human Genetics, McGill University, Montreal, Quebec, Canada

ABSTRACT

Interactive web graphics are great for communication and knowledge sharing, but are difficult to leverage during the exploratory phase of a data science workflow. Even before the web, interactive graphics helped data analysts quickly gather insight from data, discover the unexpected, and develop better model diagnostics. Although web technologies make interactive graphics more accessible, they are not designed to fit inside an exploratory data analysis (EDA) workflow where rapid iteration between data manipulation, modeling, and visualization must occur. To better facilitate exploratory web graphics that are easily distributed, we need better interfaces between statistical computing environments (e.g., the R language) and client-side web technologies. We propose the R package `animint` for rapid creation of linked and animated web graphics through a simple extension of `ggplot2`'s implementation of the Grammar of Graphics. The extension allows one to write `ggplot2` code and produce a standalone web page with multiple linked views. Supplementary material for this article is available online.

ARTICLE HISTORY

Received August 2017
Revised March 2018

KEYWORDS

Animation; Exploratory data analysis; Grammar of graphics; Multiple linked views; Statistical graphics; Web technologies

1. Introduction

For more than a half century now, statisticians have designed, built, and used interactive graphics for exploring high-dimensional data and better informing their modeling process. In fact, the ASA maintains a video library (<http://statgraphics.org/movies/>) to document and demonstrate applications of instrumental interactive statistical graphics systems such as PRIM-9 (Fisher, Friedman, and Tukey 1988), Data Viewer (Andreas et al. 1988), XGobi (Swayne, Cook, and Buja 1998), GGobi (Cook and Swayne 2007), and Mondrian (Theus 2002). These, as well as other influential systems, such as LISP-STAT (Tierney 1990) and MANET (Unwin, Hofmann, and Hawkins 1996), all have a rich support for accomplishing a wide variety of statistical analysis tasks, and most were developed before the web browser had rich graphics support.

All of these systems, as well as some more modern systems, such as `rggobi` (Duncan, Wickham, and Swayne 2016), `iplots` (Urbanek 2011), `cranvas` (Xie et al. 2013), `loon` (Waddell and Oldford 2018), etc., require a heavy set of computational dependencies to view or interact with graphics. These requirements grant the freedom to leverage libraries with sophisticated statistical functionality on-demand, but it limits the ability to share or embed such graphics in a larger document. Some of these systems allow users to create the graphics from the command-line, which as Unwin and Hofmann (2009) pointed out, allows power users to combine the strengths of a programming interface (e.g., precise, repeatable, fast, and extensible) with the strengths of a graphical interface (e.g., intuitive, forgiving, and easy-to-use). Web

technologies can certainly be used to build a similar class of system, but to capitalize on the key strengths of web technologies (e.g., accessible, portable, and composable), we must be mindful of which technologies we are requiring in such a system, and minimize those requirements whenever possible.

Generally speaking, web graphics that use purely client-side technologies (i.e., HTML, SVG, CSS, and JavaScript) are desired over client-server web applications because of their relative ease of distribution and maintenance. This is why many web-based graphing libraries like Vega (Trifacta 2014) work entirely with client-side technologies. Unfortunately, client-side technologies are not particularly well-suited for statistical computation, which we often want to leverage via dynamic controls in an interactive statistical graphics system. In this scenario, it often makes sense to introduce a client-server infrastructure to leverage functionality that is not natively supported by web browsers (e.g., R, python, etc.).

Focusing solely on the R language, there are now numerous ways to develop web applications, including the R package `shiny` (RStudio 2013), which makes it easy for R users to take their existing scripting workflow and wrap a web interface around it. Shiny is great for quickly prototyping interactive webpages that reexecute R code on-demand, but that flexibility comes at the cost of requiring a complex web server framework, which can be hard to scale, maintain, and secure sensitive information. Unfortunately, all too often, a web application framework is used to implement linked and animated graphics that could more easily be described with an idiomatic R interface which produces a purely client-side result.

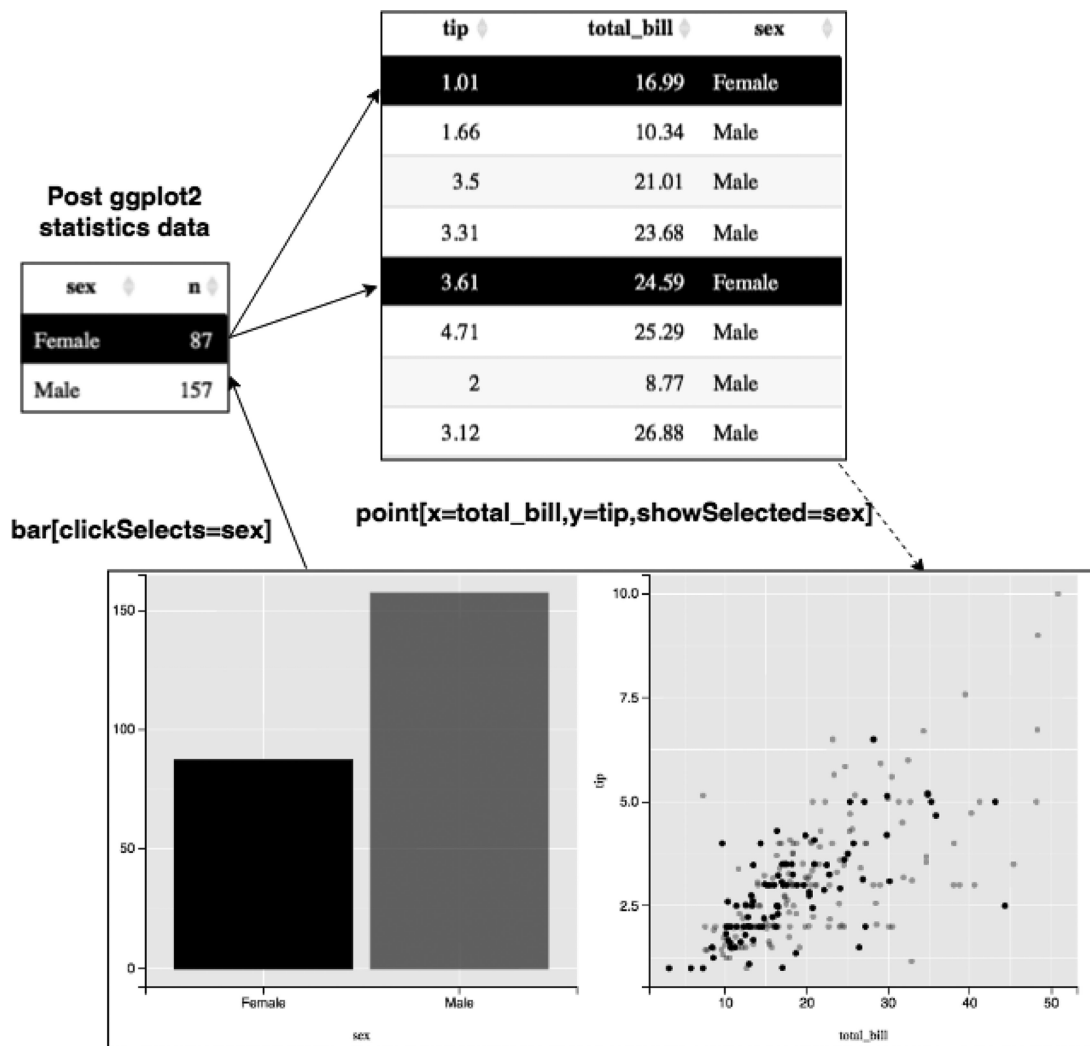


Figure 1. A graphical query of tips dataset. Left: The `clickSelects` aesthetic designates a clickable geom bar that can change a selection variable. Right: The `showSelected` aesthetic designates a geom point that responds by showing only the data which correspond to the current selection.

There are now many R packages that interface with purely client-side graphing libraries and give users the option of embedding these graphics in a larger web application. This movement has made interactive graphics a lot more accessible to newcomers and also allows power users to combine the strengths of web technologies and statistical computing. In fact, this is a large enough use case that the R package `htmlwidgets` (Vaidyanathan et al. 2018) was created to make it easier to get these interfaces to work seamlessly in any context (e.g., `shiny`, `rmarkdown`, RStudio, terminal, Jupyter notebook, etc.). In most cases, these R packages do not have great support for linking views, meaning a callback to server-side R (e.g., `shiny`) is typically required for such a task. Very recently, some `htmlwidgets` packages have gained `crosstalk` (Cheng 2016) support for linking views with purely client-side technologies, but the type of linking is purposefully restricted (e.g., 1-to-1 transient linking) since it is very difficult to standardize an API for linking arbitrary libraries. We think this is a great direction for exploratory graphics that are easily distributed, but also hope to see more opinionated approaches to this idea, like `plotly` (Sievert et al. 2018), which focus more on statistical aggregations, missing values, and selection sequences (Hofmann and Theus 1998).

We propose an extension of `ggplot2`'s layered Grammar of Graphics API to create interactive web graphics that do not require a callback to server-side R. The core idea lies in attaching metadata to graphical marks that can be used to hide/show subsets of data. The resulting framework is quite similar to what Cook, Buja, and Swayne (2007) described as brushing in multiple linked views as a database query. The assignment of data to graphical marks is done through *aesthetic mappings*, which is a term the Grammar of Graphics (Wilkinson et al. 2006) used for mapping data to visual attributes (e.g., color, shape, x , y , etc.). Typically aesthetic mappings are visual, meaning they can be easily seen in a static graphic, but our proposed aesthetic mappings control interactive properties, so they are not necessarily easily seen, but visual cues may be added to guide the user interaction. To give a small example, Figure 1 depicts a graphical query made by assigning metadata to graphical marks via the `clickSelects` and `showSelected` aesthetics. These aesthetics are essentially used to create a primary key between two tables of data, and as the name `clickSelects` suggests, queries are made by clicking directly on graphical marks, but other aesthetics could be used to support other direct manipulation events (e.g., `hoverSelects`, `brushSelects`, etc.).

Table 1. New features that `animint` adds to the grammar of graphics. These features are explained in detail starting in Section 3.1.

Feature	Type	Description
<code>clickSelects</code>	aesthetic	value(s) to select on click.
<code>showSelected</code>	aesthetic	value(s) attached to mark(s) that determine when they are shown.
<code>key</code>	aesthetic	value(s) attached to mark(s) for smooth transitions.
<code>tooltip</code>	aesthetic	information to display on hover.
<code>href</code>	aesthetic	URL link to open on click.
<code>time</code>	option	delay between animation frames.
<code>duration</code>	option	to specify smooth transitions.
<code>first</code>	option	what value(s) should be selected by default?
<code>selector.types</code>	option	should selections accumulate?
<code>selectize</code>	option	include a dropdown widget to set selection value(s) indirectly?

In addition to graphical queries, our extension supports a number of other interactive features, including animation, tooltips, and hyperlinks. A summary of these extensions and relevant additional options are provided in Table 1. There are a number of other options that can be used to control things specific to our implementation in the R package `animint` which are included with the supplementary materials.

2. Related Work

In the last section, we motivated the need for R packages that create linked interactive graphics using client-side web technologies. We also proposed an extension to `ggplot2`'s API that supports a class of graphical queries. To help further explain where our work makes contributions to the field, this section further explores related work.

It is important to acknowledge that `ggplot2` is built on top of the R package `grid`, a low-level graphics system, which is now bundled with R itself (R Core Team 2017). Neither `grid` nor base R graphics have strong support for handling user interaction, which creates a need for add-on packages. There are a number of approaches these packages take to rendering, each with their own benefits and drawbacks. Traditionally, they build on low-level R interfaces to graphical systems such as GTK+ (Lawrence and Lang 2010), Qt (Lawrence and Sarkar 2016), or Java GUI frameworks (Urbanek 2016). In general, the resulting system can be very fast and flexible, but sharing and reproducing output is usually a problem due to the heavy software requirements. Although there may be some sacrifices in performance, using the modern web browser as a rendering platform is more portable, accessible, and composable (i.e., graphics can be embedded within larger frameworks/documents).

Base R does provide a scalable vector graphics (SVG) device, `svg()`, via the Cairo graphics API (Cairo 2016). The R package `SVGAnnotation` provides functionality to post-process `svg()` output to add interactive and dynamic features (Nolan and Lang 2012). This is a powerful approach, since in theory it can work with any R graphics, but the package is self-described as a proof-of-concept which reverse engineers poorly structured `svg()` output. As a result, it is not straightforward to extend this system for linked data visualizations with advanced functionality (multiple layers, multiple plots, multiple selection variables).

The lack of well-structured SVG for R graphics motivated the `gridSVG` package which provides sensible structuring of SVG output for `grid` graphics (Murrell and Potter 2015). This package also provides some low-level tools for animating or adding interactive features, where `grid` objects must be referenced by

name. As a result, use of this interface to add interactivity to a `ggplot2` plot requires understanding of the grid naming scheme `ggplot2` uses internally. An interface where interactivity can be expressed by referencing the data to be visualized, rather than the building blocks of the graphics system, would be preferable since the former interface is decoupled from the implementation and does not require knowledge of `grid`.

In terms of the animation API, the R package `gganimate` is very similar to our system (Robinson 2016). It directly extends `ggplot2` by adding a new aesthetic, named `frame`, which splits the data into subsets (one for each unique value of the frame variable), produces a static plot for each subset, and uses the `animation` package to combine the images into a key frame animation (Xie 2013). This is quite similar but not as flexible as our system's support for animation, which we fully describe in Section 3.5. Either system has the ability to control the amount of time that a given frame is displayed, but our system can also animate the transition between frames via the `d3.transition()` API (Bostock, Oglevetsky, and Heer 2011). Smooth transitions help the animation viewer track positions between frames, which is useful in many scenarios, such as the World Bank example in Section 3.2. The `tweenr` package is similar in scope to `d3.transition()`, but operates on data frames instead of SVG elements (Pedersen 2016). One could actually use `tweenr` to implement smooth transitions in `animint`, but it would require precomputing, storing, and loading an unnecessarily large amount of data.

Smooth transitions are also useful for touring data—a suite of statistical techniques for visualization of high-dimensional data. The supplementary materials show how to implement a tour in a standalone web page via `animint` and `tourr` (Wickham et al. 2011), but it is worth noting that projections (i.e., animation frames) must be precomputed, so the functionality is limited compared to other solutions. The open-source software `GGobi` is currently the most fully featured toolkit for touring data and has support for interactive techniques such as linking, zooming, panning, and identifying (Cook and Swayne 2007). The R package `rggobi` provides an R interface to `GGobi`'s graphical interface, but it unfortunately has many software requirements. Furthermore, sharing the interactive versions of these graphics are not possible. The R package `cranvas` aims to be the successor to `GGobi`, with support for similar interactive techniques, but with a more flexible interface for describing plots inspired by the Grammar of Graphics. `Cranvas` also has many software requirements which limits its portability and accessibility.

The R package `ggvis` (Chang and Wickham 2015) is another interactive web graphics interface inspired by the Grammar of Graphics. It does not directly extend `ggplot2`,

but instead provides a brand new purely functional interface which is designed with interactive graphics in mind. It currently relies on Vega to render the SVG graphics from JSON, and the R package `shiny` to enable many of its interactive capabilities (RStudio 2013). The interface gives tremendous power to R users, as it allows one to write R functions to handle user events. This power often comes with a cost, though, as `ggvis` uses callbacks to R via `shiny` to accomplish interactivity such as linked brushing. As we outline in our supplementary materials, our system does not require server-side R, but it can also be used inside `shiny` web applications.

Another R package for interactive graphics is `iplots` (Urbanek 2011), which has several important differences compared to `animint`. Brushing of linked `iplots` is supported for single-layer plots such as scatterplots or barplots, but it is not easy to define new multi-layer interactive plots. Furthermore since `iplots` does not use the Grammar of Graphics, it is difficult to create legends and multi-panel plots. Finally, since `iplots` requires compiled C++ code for rendering on the local machine, its graphics are not as easy to share as `animint` graphics which can be viewed in a web browser.

3. Extending the Layered Grammar of Graphics

In this section, we describe in detail our extension of `ggplot2`'s layered grammar of graphics implementation (Wickham 2010). In `ggplot2`, there are five essential components that define a layer of graphical makings: data, mappings (i.e., aesthetics), geometry, statistic, and position. These simple components are easy to understand in isolation and can be combined in many ways to express a wide array of graphics. For a simple example, here is one way to create a scatterplot in `ggplot2` of variables named `<X>` and `<Y>` in `<DATA>`:

```
ggplot() + layer(
  data = <DATA>,
  mapping = aes(x = <X>, y = <Y>),
  geom = "point",
  stat = "identity",
  position = "identity"
)
```

For every geometry, `ggplot2` provides a convenient wrapper around `layer()` which provides sensible defaults for the statistic and position (in this case, both are “identity”):

```
ggplot() + geom_point(
  data = <DATA>,
  aes(x = <X>, y = <Y>)
)
```

A single `ggplot2` plot can be comprised of multiple layers, and different layers can correspond to different data. Since each graphical mark within a `ggplot2` layer corresponds to one (or

more) observations in `<DATA>`, aesthetic mappings provide a mechanism for mapping graphical selections to the original data (and vice-versa) which is essential to any interactive graphics system (Wickham et al. 2010). Thus, given a way to combine multiple `ggplot2` plots into a single view, this design can be extended to support a notion of multiple linked views, as those discussed by Ahlberg, Williamson, and Shneiderman (1991) and Buja et al. (1991).

3.1. Linking Views Via Aesthetic Mappings

Cook and Swayne (2007) used SQL queries to formalize the linked views infrastructure discussed in Ahlberg, Williamson, and Shneiderman (1991) and Buja et al. (1991). We use a similar approach to show how aesthetic mappings can be used to assign data values to graphical marks via `ggplot2` to support similar graphical queries. It is worth noting that, since these aesthetics effectively define a set of database queries that are known at print time, these queries can be made by direct manipulation of graphical marks and/or indirect manipulation via a dropdown widget, as discussed in Section 3.4. It is also worth noting that these aesthetics could be defined in such a way that they are not solely restricted to any particular direct manipulation event (e.g., mouse click), but for sake of demonstration, we restrict focus to our current `animint` implementation, which has `clickSelects` and `showSelected` aesthetics.

Consider the R code below which uses these aesthetics to create the interactive graphic depicted in Figure 1.¹ The `geom_bar()` layer in the left-hand panel is linked to the 2nd `geom_point()` layer in the right-hand panel since the `clickSelects` and `showSelected` aesthetics are mapped to a common variable, `sex`. This effectively creates a primary key relationship between the two tables used to render these graphical layers. The first `geom_point()` layer intentionally does not have a `showSelected` mapping, but has a bit of alpha transparency, so all the data are shown in light-gray, and the current selection is portrayed in black:

```
data(tips, package = "reshape2")
bar <- ggplot(tips) +
  geom_bar(aes(x = sex, clickSelects = sex))
scatter <- ggplot(tips) +
  geom_point(aes(x = total_bill, y = tip), alpha = 0.3) +
  geom_point(aes(x = total_bill, y = tip, showSelected = sex))
```

In Figure 1, there is only one selection variable which we refer to as `selected_sex`. This variable is updated whenever a bar is clicked which triggers our system to perform an SQL query of the form:

```
SELECT * FROM tips
WHERE sex IN selected_sex
```

In this example, `selected_sex` is either Male or Female (a single selected value), but as we show in later examples, a selection set can also be multiple values. Although the `clickSelects` aesthetic is tied to a mouse click event, other

¹ Interactive versions of all of the figures mentioned in this article are available with the supplementary materials, and may be viewed at <http://members.cbio.mines-paristech.fr/thocking/animint-paper-figures/>.

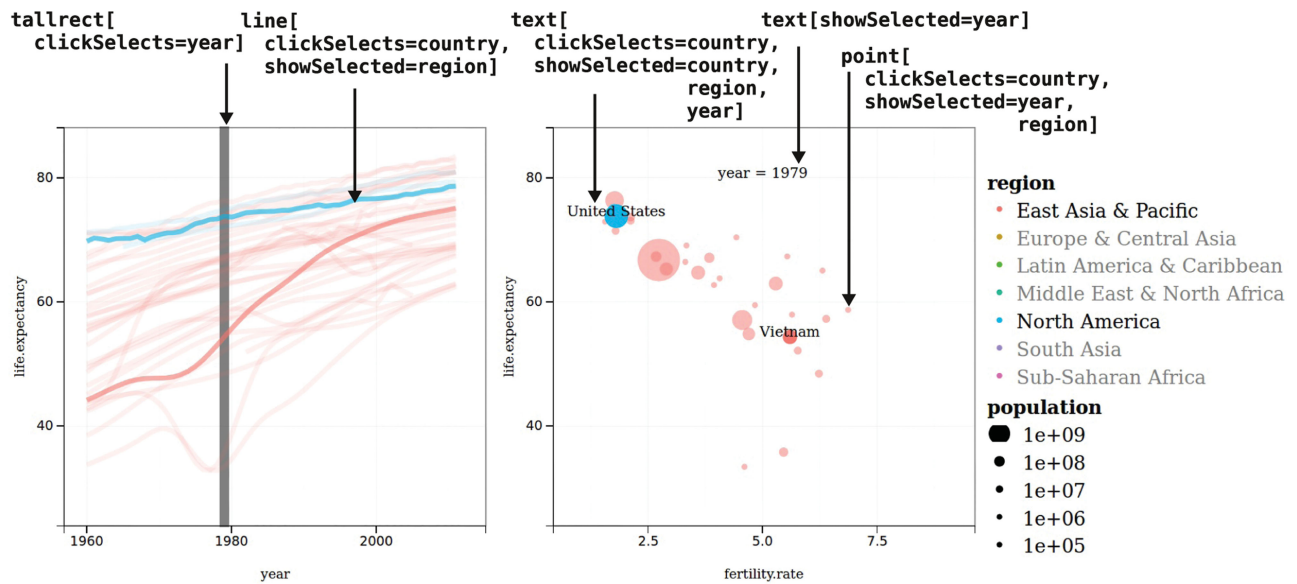


Figure 2. An interactive animation of World Bank demographic data of several countries, designed using `clickSelects` and `showSelected` aesthetics (top). Left: A multiple time series from 1960 to 2010 of life expectancy, with bold lines showing the selected countries and a vertical gray tallrect showing the selected year. Right: A scatterplot of life expectancy versus fertility rate of all countries. The legend and text elements show the current selection: year = 1979, country = {United States, Vietnam}, and region = {East Asia & Pacific, North America}.

aesthetics could easily be created to support other selection events, such as hover or click+drag. Statistically speaking, this type of interaction is useful for navigating through joint distributions conditional upon discrete values. In this sense, our extension is closely related to trellis displays (Becker, Cleveland, and Shyu 2010) and linked scatterplot brushing (Becker and Cleveland 1987). The major differences are that our conditioning is layer-specific (not plot-specific), is not tied to a particular geometry, and can be controlled through direct manipulation or animation controls.

3.2. World Bank Example

This section uses the linking framework introduced in the previous section to visualize a more complex dataset provided by the World Bank. The interactive version of Figure 2 fosters exploration of the relationship between life expectancy and fertility rate over time for 205 countries. The year 1979 and the countries United States and Vietnam are selected in the static version of Figure 2, but readers are encouraged to change the selection by clicking on the interactive version, which is provided in the supplementary materials. The interactive version also makes use of additional animation options (explained later in Section 3.5), allowing us to visualize the evolution of the relationship between life expectancy and fertility rate.

We anticipate that some `ggplot2` users will be able to reverse engineer the code which creates Figure 2, simply by looking at it. In fact, this is a big reason why `ggplot2` is so widely used: it helps minimize the amount of time required to translate an idea for a figure into computer code. Note that, in the left-hand plot of Figure 2, we have a time series of the life expectancy where each line is a country (i.e., we `group` by country) and lines are colored by region. By clicking on a line, we also want the country label to appear in the right-hand plot, so we also need to set `clickSelects=country`. Finally, by

setting `showSelected=region` and `color=region`, we can hide/show lines by clicking on the color legend entries.

```
timeSeries <- ggplot() + geom_line(
  data = WorldBank,
  aes(x = year, y = life.expectancy,
      group = country, color = region,
      clickSelects = country,
      showSelected = region)
)
```

To help point out the currently selected year, we also provide a visual cue in the form of tall rectangles to the time series plot. These tall rectangles will also serve as a way to directly modify the selected year. The tallrect geometry is a special case of a rectangle that automatically spans the entire vertical range, so we just have to specify the horizontal range via `xmin` and `xmax` aesthetics. Also, since the layered grammar of graphics allows for different data in each layer, we supply a data frame with just the unique years in the entire data for this layer:

```
years <- data.frame(year = unique(WorldBank$year))
timeSeries <- timeSeries + geom_tallrect(
  data = years,
  aes(xmin = year - 0.5, xmax = year + 0.5,
      clickSelects = year)
)
```

As for the right-hand plot in Figure 2, there are three layers: a point layer for countries, a text layer for countries, and a text layer to display the selected year. By clicking on a point, we want to display the country text label and highlight the corresponding time series on the left-hand plot, so

we set `clickSelects=country` in this layer. Furthermore, we only want to show the points for the selected year and region, so we also need `showSelected=year` and `showSelected2=region`.

```
scatterPlot <- ggplot() + geom_point(
  data = WorldBank,
  aes(x = fertility.rate, y = life.expectancy,
      color = region, size = population,
      clickSelects = country,
      showSelected = year,
      showSelected2 = region)
)
```

Note that any aesthetics containing the substring `showSelected` (including `showSelected2`) are interpreted as `showSelected` variables, and combined together using the intersection operation. In the example above, that means that a point will be drawn for the currently selected combination of year and region, as in the following SQL query,

```
SELECT * FROM WorldBank
WHERE year IN selected_year
AND region IN selected_region
```

Below, the text layer for annotating selected countries is essentially the same as the point layer, except we assign the country name to the `label` aesthetic:

```
scatterPlot <- scatterPlot + geom_text(
  data = WorldBank,
  aes(x = fertility.rate, y = life.expectancy,
      label = country,
      showSelected = country,
      showSelected2 = year,
      showSelected3 = region)
)
```

Lastly, to help identify the selected year when viewing the scatterplot, we add another layer of text at a fixed location.

```
scatterPlot <- scatterPlot + geom_text(
  data = years, x = 5, y = 80,
  aes(label = paste("year =", year),
      showSelected = year)
)
```

In summary, this section shows an example of how the proposed `clickSelects` and `showSelected` aesthetics can be used with several different geoms (line, point, text, `tallrect`), each of which can potentially display a different dataset. In each case, we use `clickSelects` to declare a geom that when clicked updates the current selection, and we use `showSelected` to declare a geom which responds to such changes

by updating the set of displayed data. In the next sections, we further explore options that allow us to accumulate selections, update selections indirectly, and automate selection updates.

3.3. Linking and Multiple Selection

Linking is declared in R code by putting ggplots with common `clickSelects` and `showSelected` aesthetics together in a list. For example, we can link the ggplots from the previous section by including them together in the following list:

```
viz <- list(
  timeSeries = timeSeries,
  scatterPlot = scatterPlot
)
```

Linking is accomplished because the two ggplots declared `clickSelects` and `showSelected` aesthetics that refer to common variable names (`region`, `year`, `country`). For each such selection variable, our system updates the set of selected values in response to mouse clicks on `clickSelects` geoms, and then updates the corresponding data which is displayed for `showSelected` geoms.

Note that the `viz` list above can also contain numerous options which are listed in [Table 1](#). For example, the `selector.types` option controls whether or not selections for a given variable accumulate (single or multiple selected values). This sort of logic has also been interpreted as transient versus persistent selection (Cook and Swayne 2007).

```
viz$selector.types <- list(
  year = "single",
  country = "multiple",
  region = "multiple"
)
```

The code above declares `year` as a single selection variable, which means that only a single year may be selected at a time (clicking a geom with `clickSelects=year` will change the selection to the corresponding year). The `country` and `region` variables are declared as multiple selection variables, which can have multiple selected values at a time (clicking a geom with `clickSelects=country` will add/remove that country to/from the selection set).

3.4. Direct Versus Indirect Manipulation

Graphical queries via direct manipulation require direct user interaction with graphical elements, but it is not necessarily easy to find value(s) of interest in the graphical space. For this reason, `animint` also provides dropdown widgets for executing graphical queries via indirect manipulation. For example, when viewing the interactive version of [Figure 2](#), suppose our goal is to compare the United States to Thailand. Direct manipulation is not very useful in this case since it is not necessarily easy to identify and select these countries based solely on the graphics.

Toggle selected value

year

region

country

Figure 3. Using dropdown widget(s) to execute graphical queries via indirect manipulation. This example shows how one could search for and highlight Thailand in Figure 2.

Figure 3 shows what the user sees after typing “th” in the search box. Note that these dropdowns support selection of multiple values and are coordinated with selections made via direct manipulation.

3.5. Animation and Smooth Transitions

Animation is declared using the `time` option, which specifies a selection variable that will be automatically updated over time, as well as a time delay in milliseconds. The code below declares the `year` variable to be animated every 3 sec:

```
viz$time <- list(variable = "year", ms = 3000)
```

Animation is useful in the World Bank data visualization because it shows how the bi-variate relationship between fertility rate and life expectancy changes over time. Animation clearly shows how many countries progress from low life expectancy and high fertility rate in early years, to high life expectancy and low fertility rate in later years.

Finally, the `duration` option specifies the amount of time used to smoothly transition between selections (with linear easing). Smooth transitions help the viewer track geoms before and after an update to the selection set. For example, in the code below we declare a 1 sec smooth transition on the `year` variable, to more easily track the points on the scatterplot:

```
viz$duration <- list(year = 1000)
```

Note that for accurate interpretation of smooth transitions, the new key aesthetic must be specified. The key aesthetic is used to match data elements before and after the smooth transition. In the World Bank example, we would need to specify `aes(key=country)` for the points and text in the scatterplot.

3.6. Storing and Restoring state

When sharing an interactive visualization with others, it can often be helpful to share interesting state(s) of the visualization. In `animint`, states can be serialized in a URL link and/or specified at the command line via the `first` option. The code below declares that the first selection of the `country` variable is the set of two countries, United States and Vietnam:

```
viz$first <- list(country = c("United States", "Vietnam"))
```

3.7. Compiling and Rendering

Supplying the `viz` list of ggplots and rendering options to the `animint2dir()` function will save all the files necessary for rendering the visualization:

```
animint::animint2dir(viz)
```

As shown in supplementary Figure 1, the `animint` system consists of two parts: the compiler and the renderer. The compiler is R code that converts a list of ggplots and options to a JSON plot meta-data file and a tab-separated values (TSV) file database. The renderer consists of HTML and JavaScript files, which can be easily hosted along with the TSV and JSON files on any web server. The interactive plots can be viewed by opening the `index.html` page in any modern web browser. Note that `animint` currently depends on a fork of `ggplot2` (<https://github.com/faizan-khan-iit/ggplot2/tree/validate-params>) that contains some minor modifications which are needed to support interactive rendering on web pages. Additional implementation details are available in the supplementary materials.

4. Exploring Scope with Examples

This section attempts to demonstrate the range of visualizations that are possible with `animint`. In particular because of its support for interaction and animation, it excels at display of interactive maps with time-varying data. We give two such examples below. A handful of other examples are provided with the supplementary materials.

4.1. Tornadoes in the United States

One of the strong points of the system we propose is display of multi-layer plots such as maps with time-varying data. For example, Figure 4 shows a visualization of U.S. tornado data from 1950 to 2012. This data visualization consists of two multi-layer plots with two interaction variables, `year` and `state`.

The left plot is a map which shows state borders using a polygon with `clickSelects=state`. The currently selected state is shown using semitransparency, and other states can be selected by clicking them. The state map plot uses geoms with `showSelected=year` to show tornado paths (segment geom) and endpoints (point geom) for the currently selected year (which is emphasized with a text geom above the map).

The right plot uses several geoms to show details for the currently selected state and year. A bar geom shows a time series of tornado counts for the selected state (`showSelected=state`), which can be clicked to change the currently selected year (`clickSelects=year`). A text geom at the top of the plot shows the currently selected state (`showSelected=state`), and a text geom at the bottom emphasizes the tornado count for the selected year (using `showSelected` variables for both `state` and `year`).

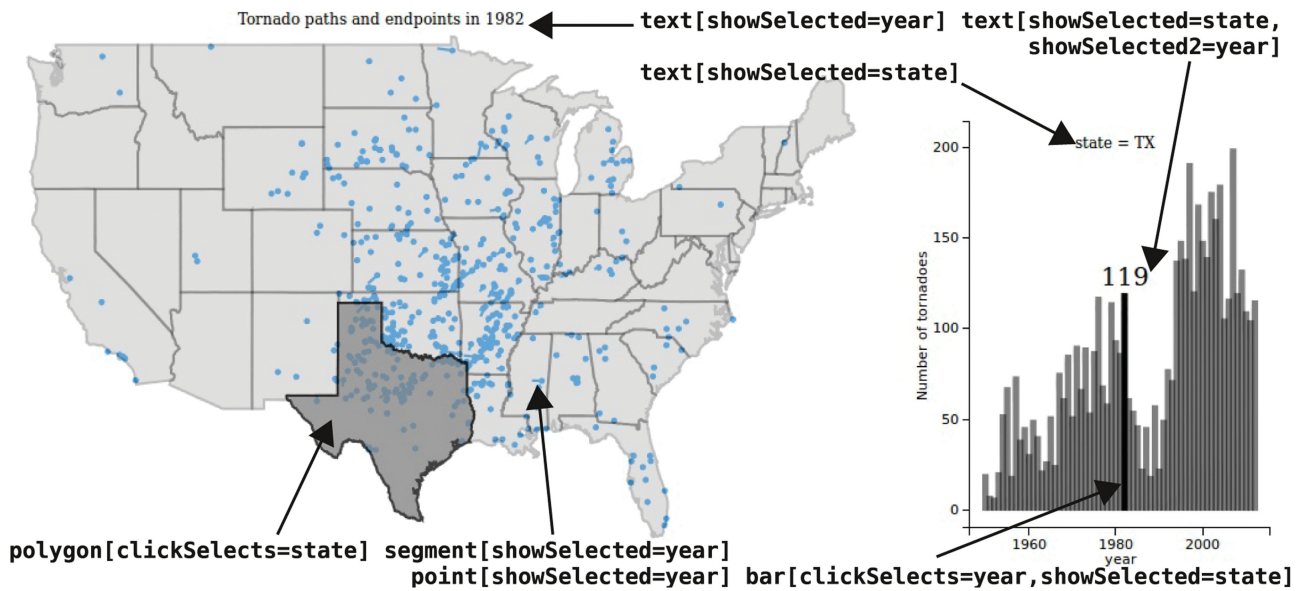


Figure 4. Interactive animation of U.S. tornadoes from 1950 to 2012. This figure depicts a scenario where the user queried Texas (by clicking the map), and the year 1982 (by clicking the bar chart). In addition to the graphical elements being highlighted as a visual clue of what query is being made, this visualization includes dynamic text labels reflecting the query.

These interactions can be useful for discovering patterns in the data, and for suggesting models that can describe or predict tornado paths.

4.2. Central American Climate Data

A more complex map data visualization example is shown in Figure 5, which depicts climate time series data observed in Central America. There are two interaction variables, time and region.

Two maps in the upper left display borders of the countries in and near Central America. Unlike the previous example with U.S. states, the country borders are static (clicking has no effect). For the currently selected time, rect geoms with `showSelected=time` show the spatial distribution of sea surface temperature as well as its deviation from the monthly norm. Since `clickSelects=region` is specified, clicking a rect changes the currently selected region, which is emphasized with a black border. These plots facilitate visualization of the spatial distribution of the climate variables, and how they change over time.

The plots below the maps use lines to show time series of the climate variables. Since `clickSelects=region` is specified, clicking a line changes the currently selected region, which is emphasized with a purple color. A semitransparent tallrect shows the currently selected time; other tallrects can be clicked to update the time (`clickSelects=time`). These plots make it easy to select different times and regions, and to make comparisons between times and regions.

Scatterplots on the right use `showSelected` variables with point and text geoms, to show the joint distribution of the two temperature variables for the selected time (top) and region (bottom). The plots use `clickSelects` to emphasize the currently selected region (top) and time (bottom), and are useful for visualizing normality and outliers in the joint distribution.

5. Limitations and Future Work

Our implementation of the `ggplot2` extension proposed in Section 3.1 has a number of limitations. Most notably, an interactive statistical graphics system should be able to dynamically compute statistical aggregations based on new user input (e.g., compute a new linear model based on a set of newly brushed points). In theory, the extension can support dynamic statistical aggregations specified via a `ggplot2` layer, but it is not yet clear to us how one would translate every possible R function to JavaScript, so this is not currently implemented in `animint`. It may be worthwhile exploring compilation of R functions using a technology like `WebAssembly` so that the web browser can run them without an external web server running R. Nevertheless, it is currently possible to workaround this problem somewhat by precomputing every possible aggregation ahead of time. (If the total number of selection states is fairly small, this approach works, but it does not scale very well. For an example of precomputing states and exploring those limitations, see the supplementary materials.)

Numerous other limitations in our current implementation derive from the fact that some plot features are computed once during the compilation step, and remain static on a rendered plot. For example, users are not able to dynamically alter variable mappings, transformations, or axis scaling. Most of these limitations can be resolved by adding interactive widgets to recompile plot(s) via a callback to R. For this reason, `animint` makes it easy to embed visualizations inside of `shiny` web applications—refer to the supplementary materials for an example.

Other limitations could also be addressed by adding a few other aesthetics or options to the list provided in Table 1. More specifically, one could add support for more forms of direct manipulation by adding `hoverSelects` and `brushSelects` aesthetics, option(s) for more control over the selection styling (e.g., color, opacity, etc.), and options for providing

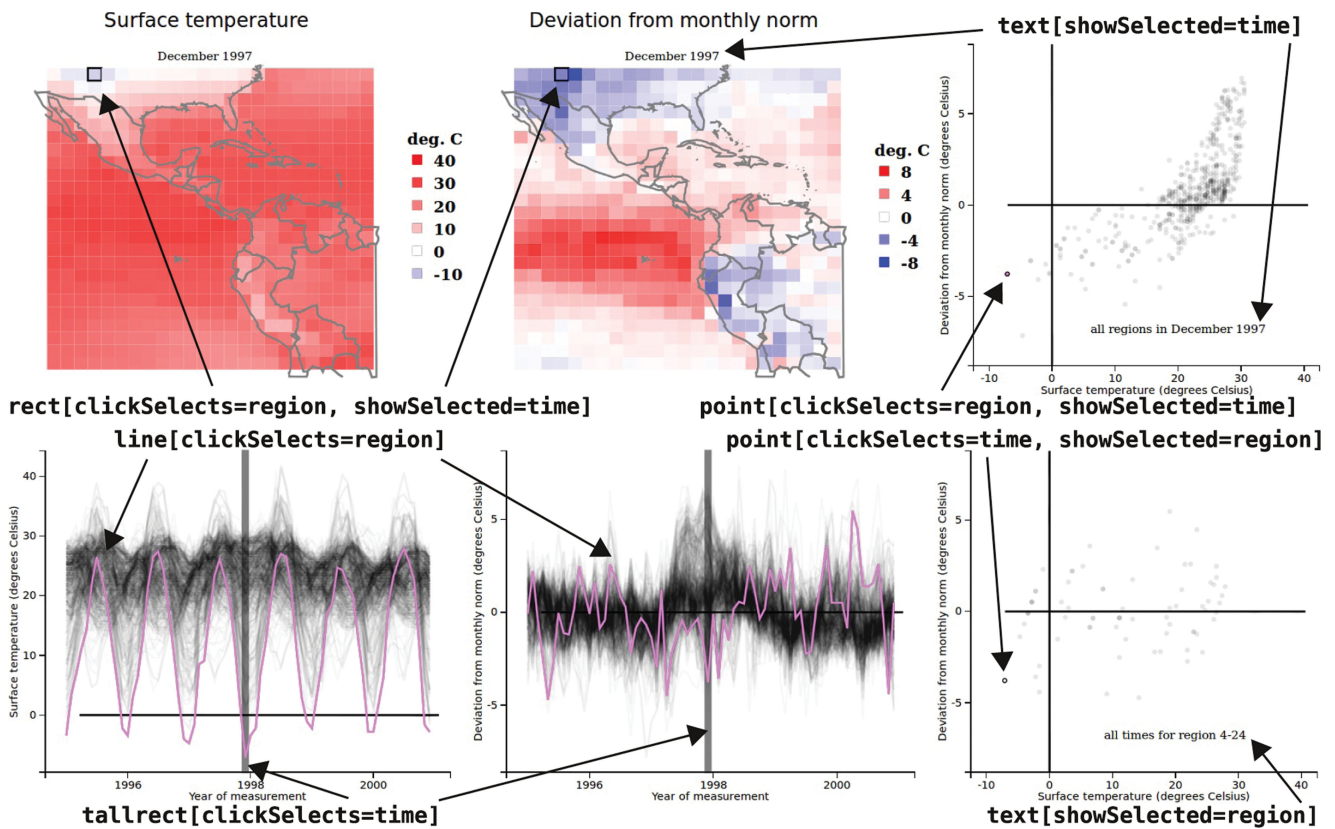


Figure 5. Visualization containing six linked, interactive, animated plots of Central American climate data. Top: For the selected time (December 1997), maps displaying the spatial distribution of two temperature variables, and a scatterplot of these two variables. The selected region is displayed with a black outline, and can be changed by clicking a rect on the map or a point on the scatterplot. Bottom: Time series of the two temperature variables with the selected region shown in violet, and a scatterplot of all times for that region. The selected time can be changed by clicking a background tallrect on a time series or a point on the scatterplot. The selected region can be changed by clicking a line on a time series.

more visual cues for graphical objects that trigger graphical queries (currently, hovering on such objects will change their transparency). There are other types of interaction that could be added without adding to the extension at all (e.g., zooming, panning, and plot resizing).

6. Conclusion

Interactive graphics can augment data exploration, and lead to better understanding by allowing one to quickly answer follow-up questions; but to be practically useful, one should be able to iterate quickly and share easily. Interactive statistical graphics have traditionally had heavy software requirements since it is common to dynamically execute statistical aggregations based on user input. However, if we wish to bring interactive statistical graphics to the web in a responsible way, we should explore how we can enable common tasks such as animation and linked views without a complex client-server infrastructure. Our simple extension of `ggplot2`'s layered grammar of graphics enables a set of common interactive tasks (e.g., graphical queries in multiple views, animation, tooltips, hyperlinks, etc.) without requiring a complex client-server infrastructure.

Supplementary Materials

Interactive Figures and Reproducible Research Statement

The source code to create this article and its figures is online at <https://github.com/tdhock/animint-paper/> and the interactive figures can be

viewed at <http://members.cbio.mines-paristech.fr/thocking/animint-paper-figures/>.

Acknowledgments

The authors thank animint users MC Du Plessis, Song Liu, Nikoleta Juretic, and Eric Audemard who have contributed constructive criticism and helped its development.

ORCID

Carson Sievert  <http://orcid.org/0000-0002-4958-2844>

References

Ahlberg, C., Williamson, C., and Shneiderman, B. (1991), "Dynamic Queries for Information Exploration: An Implementation and Evaluation," in *ACM Chi '92 Conference Proceedings* (Vol. 21), ACM, pp. 619–626. [302]

Andreas, B., Hurley, C., Asimov, D., and McDonald, J. A. (1988), "Elements of a Viewing Pipeline for Data Analysis," in *Dynamic Graphics for Statistics*, eds. William S. Cleveland and Marylyn E. McGill, Belmont, CA: Wadsworth, Inc. [299]

Becker, R. A., and Cleveland, W. S. (1987), "Brushing Scatterplots," *Technometrics*, 29, 127–142. [303]

Becker, R. A., Cleveland, W. S., and Shyu, M.-J. (2010), "The Visual Design and Control of Trellis Displays," *Journal of Computational and Graphical Statistics*, 19, 3–28. [303]

Bostock, M., Oglevetsky, V., and Heer, J. (2011), "D3 Data-Driven Documents," *IEEE Transactions on Visualization and Computer Graphics*, 17, 2301–2309. [301]

- Buja, A., McDonald, J. A., Michalak, J., and Stuetzle, W. (1991), "Interactive Data Visualization Using Focusing and Linking," *IEEE Proceedings of Visualization*, February, 1–8. [302]
- Cairo (2016), "Cairo: A Vector Graphics Library," available at <http://cairographics.org/> [301]
- Chang, W., and Wickham, H. (2015), ggvis: Interactive Grammar of Graphics, available at <https://CRAN.R-project.org/package=ggvis> [301]
- Cheng, J. (2016), "Crosstalk: Inter-Widget Interactivity for HTML Widgets," available at <https://CRAN.R-project.org/package=crosstalk> [300]
- Cook, D., Buja, A., and Swayne, D. F. (2007), "Interactive High-Dimensional Data Visualization," *Journal of Computational and Graphical Statistics*, December, 1–23. [300]
- Cook, D., and Swayne, D. F. (2007), *Interactive and Dynamic Graphics for Data Analysis: With R and GGobi. Use R!* New York: Springer, available at <http://www.ggobi.org/book/> [299,301,302,304]
- Duncan, T. L., Wickham, H., and Swayne, D. (2016), "Interface Between R and GGobi." [299]
- FisherKeller, M. A., Friedman, J. H., and Tukey, J. W. (1988), "PRIM-9, An Interactive Multidimensional Data Display and Analysis System," in *Dynamic Graphics for Statistics*, pp. 91–109. [299]
- Hofmann, H., and Theus, M. (1998), "Selection Sequences in MANET," *Computational Statistics*, January, 1–12. [300]
- Lawrence, M., and Lang, D. T. (2010), "RGtk2: A Graphical User Interface Toolkit for R," *Journal of Statistical Software*, 37, 1–52. available at <http://www.jstatsoft.org/v37/i08/> [301]
- Lawrence, M., and Sarkar, D. (2016), "Interface Between R and Qt," available at <https://github.com/ggobi/qtbase> [301]
- Murrell, P., and Potter, S. (2015), "gridSVG: Export 'grid' Graphics as SVG," available at <https://CRAN.R-project.org/package=gridSVG> [301]
- Nolan, D., and Lang, D. T. (2012), "Interactive and Animated Scalable Vector Graphics and R Data Displays," *Journal of Statistical Software*, 46, 1–88. available at <http://www.jstatsoft.org/v46/i01/> [301]
- Pedersen, T. L. (2016), "tweenr: Interpolate Data for Smooth Animations," available at <https://CRAN.R-project.org/package=tweenr> [301]
- R Core Team (2017), *R: A Language and Environment for Statistical Computing*, Vienna, Austria: R Foundation for Statistical Computing. available at <http://www.R-project.org/> [301]
- Robinson, D. (2016), "gganimate: Create Easy Animations with ggplot2," available at <http://github.com/dgrtwo/gganimate> [301]
- RStudio (2013), "Shiny: Easy Web Applications in R," available at <http://www.rstudio.com/shiny/> [299,302]
- Sievert, C., Parmer, C., Hocking, T., Chamberlain, S., Ram, K., Corvellec, M., and Despouy, P. (2018), *Plotly: Create Interactive Web Graphics via 'Plotly.js'*. [300]
- Swayne, D. F., Cook, D., and Buja, A. (1998), "XGobi: Interactive Dynamic Data Visualization in the X Window System," *Journal of Computational and Graphical Statistics*, 7, 113–130. [299]
- Theus, M. (2002), "Interactive Data Visualization Using Mondrian," *Journal of Statistical Software*, 7, 1–9. Available at <http://www.jstatsoft.org/v07/i11>. [299]
- Tierney, L. (1990), *LISP-Stat: An Object Oriented Environment for Statistical Computing and Dynamic Graphics*, New York: Wiley-Interscience. [299]
- Trifacta (2014), "Vega: A Declarative Visualization Grammar," available at <http://trifacta.github.io/vega/> [299]
- Unwin, A., and Hofmann, H. (2009), "GUI and Command-line - Conflict or Synergy?" *Proceedings of the St Symposium on the Interface*, September, 1–11. [299]
- Unwin, A., Hofmann, H., and Hawkins, G. (1996), "Interactive Graphics for Data Sets with Missing Values - Manet," *Journal of Computational and Graphical Statistics*, 4. [299]
- Urbanek, S. (2011), "IPlots eXtreme: Next-Generation Interactive Graphics Design and Implementation of Modern Interactive Graphics," *Computational Statistics*, 26, 381–393. available at <https://doi.org/10.1007/s00180-011-0240-x> [299,302]
- (2016), "rJava: Low-level R to Java Interface," available at <https://CRAN.R-project.org/package=rJava> [301]
- Vaidyanathan, R., Xie, Y., Allaire, J. J., Cheng, J., and Russell, K. (2018), "htmlwidgets: HTML Widgets for R," available at <https://github.com/rmnnathv/htmlwidgets> [300]
- Waddell, A., and Oldford, R. W. (2018), "loon: Interactive Statistical Data Visualization," available at <https://CRAN.R-project.org/package=loon> [299]
- Wickham, H. (2010), "A Layered Grammar of Graphics," *Journal of Computational and Graphical Statistics*, 19, 3–28. [302]
- Wickham, H., Cook, D., Hofmann, H., and Buja, A. (2011), "Tourr: An R Package for Exploring Multivariate Data with Projections," *Journal of Statistical Software*, 40, 1–18. [301]
- Wickham, H., Lawrence, M., Cook, D., Buja, A., Hofmann, H., and Swayne, D. F. (2010), "The Plumbing of Interactive Graphics," *Computational Statistics*, April, 1–7. [302]
- Wilkinson, L., Wills, D., Rope, D., Norton, A., and Dubbs, R. (2006), *The Grammar of Graphics*, New York: Springer. [300]
- Xie, Y. (2013), "Animation: An R Package for Creating Animations and Demonstrating Statistical Methods," *Journal of Statistical Software*, 53, 1–27. available at <http://www.jstatsoft.org/v53/i01/> [299]
- Xie, Y., Hofmann, H., Cook, D., Cheng, X., Schloerke, B., Vendettuoli, M., Yin, T., Wickham, H., and Lawrence, M. (2013), *Interactive Statistical Graphics Based on Qt*. [301]