




Manipulable Semantic Components: A Computational Representation of Data Visualization Scenes

Zhicheng Liu , Chen Chen , and John Hooker 

Abstract—Various data visualization applications such as reverse engineering and interactive authoring require a vocabulary that describes the structure of visualization scenes and the procedure to manipulate them. A few scene abstractions have been proposed, but they are restricted to specific applications for a limited set of visualization types. A unified and expressive model of data visualization scenes for different applications has been missing. To fill this gap, we present Manipulable Semantic Components (MSC), a computational representation of data visualization scenes, to support applications in scene understanding and augmentation. MSC consists of two parts: a unified object model describing the structure of a visualization scene in terms of semantic components, and a set of operations to generate and modify the scene components. We demonstrate the benefits of MSC in three applications: visualization authoring, visualization deconstruction and reuse, and animation specification.

Index Terms—data visualization, scene abstraction, visualization model

1 INTRODUCTION

To create and render a data visualization, we must first assemble a scene—a data structure that includes all the visual elements and components, their attributes, relationships, and constraints in the visualization. A visualization scene can be manually assembled, for example, by using an XML-based markup language based on the SVG (Scalable Vector Graphics) specification [1], or by drawing shapes on a canvas in a vector graphics editor. Since such processes are tedious for complex visualization designs, two popular approaches have been proposed to facilitate expressive scene assembly: scene manipulation libraries and declarative specifications. An example of a scene manipulation library is D3 [8], which uses the DOM (Document Object Model) to represent a visualization scene, and provides functions for programmers to select and modify scene elements. Declarative specifications, on the other hand, decouple the description of desired visualization from scene assembly. Using languages such as ggplot2 [51] and Vega-Lite [42], programmers only need to specify their intents in declarative specifications, which are compiled into visualization scenes.

Both approaches have been successful in facilitating scene assembly. However, recent developments in various visualization applications postulate the need for going beyond scene assembly to support scene understanding and augmentation. For instance, visualization reverse-engineering involves an automated system deconstructing a scene to understand the semantic roles of visual objects and their relationships. Another example is interactive visualization authoring, where users can interactively edit [6], reuse [11, 13], and annotate [36] static visualizations, as well as define interactive [44] and animated behaviors [15, 50] through a graphical user interface.

Existing scene manipulation libraries and declarative specifications do not provide vocabularies at an appropriate level for such applications. Libraries like D3 use SVG as the scene graph model, which does not capture high-level semantic abstractions of visualization scenes [10, 15]. Declarative specifications, by decoupling specification from scene assembly, hide details on how a scene is represented (e.g., the hierarchical organization of the visual objects) and manipulated (e.g., updating the number of marks based on the values of a data variable). They do not provide native support for directly referencing, selecting, and editing scene objects. Researchers thus resort to defining their own scene

abstractions. For example, Chartreuse identified two primary mark structures (units and sequences) and six update operations (morph, move, fix, recolor, repeat, and partition) to reverse-engineer and reuse infographics bar charts [13]; Canis [15] defines a new scene model called dSVG to describe the data-enriched semantic components in SVG charts for animation authoring. Such efforts, however, are restricted to specific applications for a limited set of visualizations. A unified and expressive model of data visualization scenes for different applications has been missing.

To fill this gap, we present Manipulable Semantic Components (MSC), a computational representation of data visualization scenes. By computational representation, we refer to the format and methods that define how a data visualization scene is represented and manipulated within a computational system. MSC consists of two parts: a unified *object model* describing the structure of a visualization scene in terms of semantic components, and a set of *operations* to generate and modify the scene components.

In the object model, we identify the following major types of semantic components: visual elements (e.g., mark, glyph, collection), data scope, encodings (i.e., mapping between data attribute and visual channel), algorithmic layouts (e.g., grid, stack, packing), relational constraints (e.g., alignment) and view configuration. Each of these components has a set of associated operations. Specifically, we identify *generative* operations that create or remove visual elements based on data (e.g., repeat, divide, densify, classify, repopulate, and stratify), and *modificative* operations that update visual properties of elements or relationships between elements (e.g., apply visual encoding, customize scale, apply layout, apply constraint).

We demonstrate the benefits of MSC and how it supports different applications in three case studies: visualization authoring, visualization deconstruction and reuse, and animation authoring.

2 MANIPULABLE SEMANTIC COMPONENTS IN A DIVERGING STACKED BAR CHART: AN EXAMPLE

For an example of manipulable semantic components, consider the diverging stacked bar chart in Fig. 1a. This chart visualizes a hypothetical dataset reporting people's opinions on a subject matter, broken down by age (below 30, 30 - 50, 50 - 70, above 70) and responses (strongly agree, agree, disagree, strongly disagree). The underlying data table consists of three columns: age, response, and pct (percentage of people with a particular response within an age group). The scene of this chart consists of the following semantic components (Fig. 1b): **marks** such as rectangles and texts are the graphical primitives; **groups** organize the rectangle marks by the age values; **layouts** specify the spatial relationships of rectangle marks within and across groups; **encodings** bind data attributes (e.g., pct) to visual channels (e.g., rectangle width); **constraints** enforce inter-object spatial relationships (e.g., light blue rectangles are aligned to the right, text marks are affixed to the rectangle

• Zhicheng Liu, Chen Chen, and John Hooker are with University of Maryland, College Park. E-mail: {leozcli, cchen24}@umd.edu, jhooker@terpmail.umd.edu

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

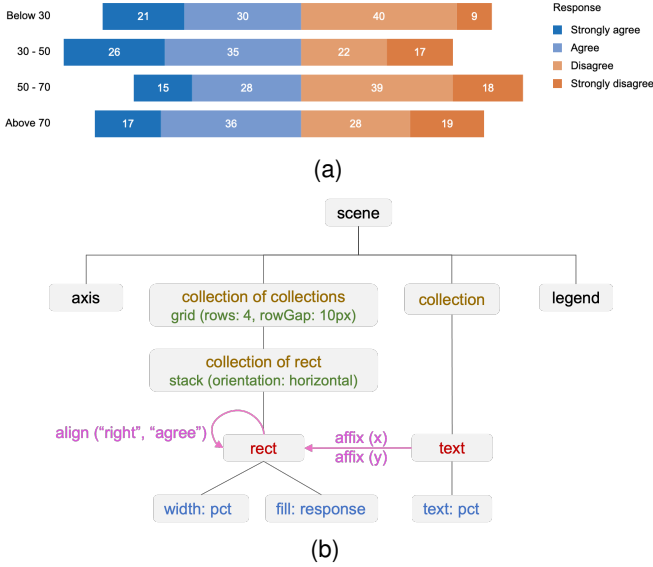


Fig. 1: Manipulable semantic components in a diverging stacked bar chart. The scene abstraction not only captures hierarchical relationships between objects but also relationships like constraints and encodings.

marks at the center). A component can have *parameters*, e.g., a stack layout component has orientation and gap parameters.

In addition to describing the scene structure, MSC also describes the generative procedure to create and modify these components. Figure 2 shows major steps of creating this chart using the MSC operations:

- create a rectangle mark;
- repeat the rectangle by age, resulting in a collection of four rectangles in a 4×1 grid layout, each representing an age group;
- divide each rectangle by response, splitting it into four rectangles in a horizontal stack layout, each representing a response type;
- encode the rectangles' width by pct and fill color by response;
- align the light blue rectangles to the right to better show the divergence of opinions;
- repeat an initial text item by pct to create the labels, and affix the position of the texts to the center of the rectangles.

During this process, we can also freely *resize* or *translate* the objects. The encodings and constraints are enforced to ensure the design represents data faithfully.

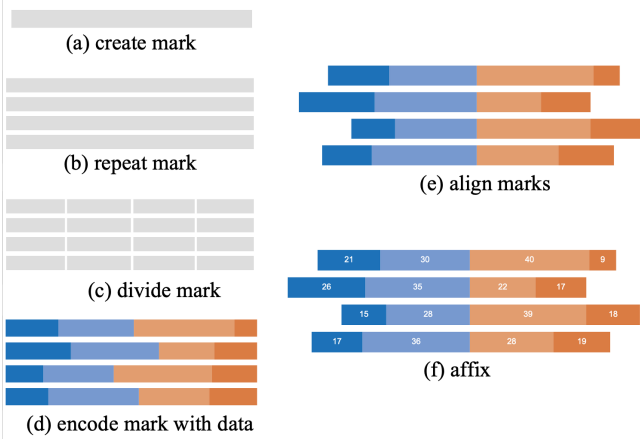


Fig. 2: Creating a diverging bar chart using the MSC operations.

3 OBJECT MODEL: SEMANTIC COMPONENTS

In MSC, a *scene* consists of semantic components including visual elements, layouts, encodings, constraints, and view configuration. MSC focuses on two kinds of *datasets*: tables and networks [30]. A *table*

consists of a set of data *items* (i.e., rows), and each item contains values for a set of *attributes* (i.e., columns). A *network* also contains items, in addition, it contains *links* that connect the items. A tree is a special type of network and thus covered in MSC. A dataset can be used across multiple scenes, and a scene can use multiple datasets.

Given an initialized scene and a dataset, we can create or modify a visualization through the following semantic components: marks, data scope, groups, auxiliary visual elements, visual encodings and scales, algorithmic layouts, relational constraints, and view configuration.

3.1 Primary Visual Elements: Marks, Vertices & Segments

Marks serve as the fundamental building blocks of data visualization. Fig. 3 shows the fourteen mark types supported in MSC, as well as how certain types of marks can be derived by applying appropriate operations to some primitive marks (detailed in Sec. 4). Marks can represent either items or links [30]. For instance, a line mark can either represent an item in a slope graph, or a link in a node-link diagram.

Definition 1 (Mark Type). The type of a mark m is defined as: $type(m) \in \{rectangle, circle, line, path, text, image, band, area, ring, pie, polyline, arc, polygon, geoPolygon\}$

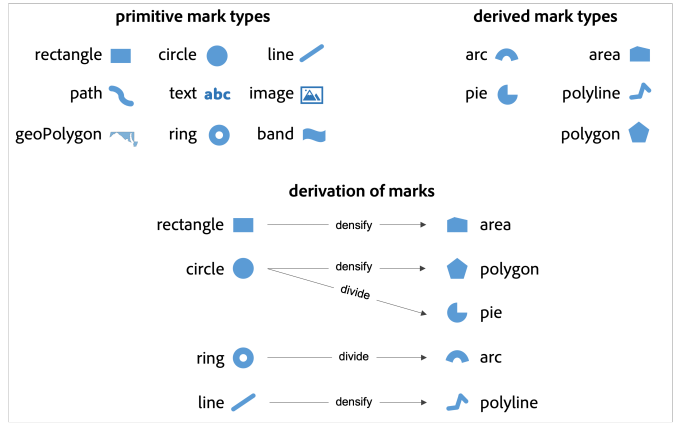


Fig. 3: Marks supported in MSC and derivation of marks.

Marks in the form of geometric shapes are composed of *vertices* and *segments*. A rectangle, for example, consists of four vertices connected by four line segments; a polyline consists of multiple vertices connected by line or curve segments.

Visual Channels. Each type of mark has a set of associated *visual channels* (e.g., width, height for rectangle, radius for circle) that determine the mark's visual appearance. Vertices and segments in a mark also have a set of associated visual channels (e.g., x & y positions of vertices in an area mark, y position of the top and bottom segments in a rectangle mark, and stroke width of segments in a polyline).

3.2 Data Scope

A mark represents one or more data items from a dataset. For example, each rectangle mark in Fig. 1a represents a distinct data item with a unique combination of age, response, and pct values. We define the data scope of a visual element e to be the data item(s) represented by the element: $e.data$.

Attribute Value and Aggregator. We use $e.data[\alpha]$ to denote the value(s) of attribute α in element e 's data scope. The element can be a single mark or a group of marks. For instance, if e refers to the bottom left blue bar in Fig. 1a, then $e.data["pct"] = 17$; if e refers to the group of four bars in the bottom row in Fig. 1(f), then $e.data["pct"] = [17, 36, 28, 29]$. In the latter case, since the attribute ("pct") is quantitative, we can define an *aggregator* (e.g., max, min, count, sum) to compute a single value: $max(e.data["pct"]) = 36$, and $min(e.data["pct"]) = 17$. For simplicity, we use $e.data[\alpha]$ in the rest of this paper to refer to an element's scope attribute value, irrespective of the number of data items in the data scope.

3.3 Primary Visual Elements: Groups

MSC defines three types of groups: glyph, collection, and composite. Like marks, groups are visual elements with associated visual channels (commonly x and y positions) and can have data scopes.

Glyph. Multiple marks can be grouped to create a glyph. Fig. 4 shows example glyphs in four different visualizations.

Definition 2 (Glyph). A glyph has the following properties:

1. All the marks in a glyph have the same data scope.
2. The data scope of a glyph is the same as the data scope of any of its member marks.

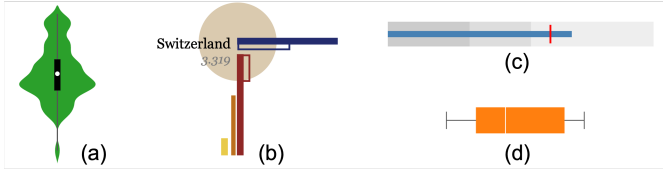


Fig. 4: Example glyphs in (a) a violin plot, (b) the “Brain drain” visualization [4], (c) a bullet chart, and (d) a box-and-whisker plot.

Definition 3 (Glyph Type). Given two glyphs G_i and G_j , they have the same type (i.e., $\text{type}(G_i) = \text{type}(G_j)$) if the following conditions are met:

1. G_i and G_j have the same number of constituent marks.
2. for every mark m_a in G_i , there exists a unique mark m_b in G_j where $\text{type}(m_a) = \text{type}(m_b)$, and vice versa.

Collection. The main content of a data visualization is typically a collection of glyphs or marks (Fig. 5). No two members in a collection can share the same data scope. For example, each rectangle mark in Fig. 5(a) represents a data item with a unique month value. Collections can be nested: the members of a collection are collections. In Fig. 1f, we have four collections of stacked bars, and they are the members of a larger collection that corresponds to the main chart content.

Definition 4 (Collection). A collection of visual elements has the following properties:

1. the members in a collection must be all marks, all glyphs, or all collections.
2. any two members e_i and e_j in a collection must have the same type (per Definitions 1, 3, and 5): $\text{type}(e_i) = \text{type}(e_j)$.
3. the data scopes of all the members in the same collection share the same attributes.
4. the data scopes of any two members e_i and e_j in a collection do not overlap: $e_i.\text{data} \cap e_j.\text{data} = \emptyset$.
5. the data scope of a collection C is the union of the data scopes of its members: $C.\text{data} = \bigcup_{e_i \in C} e_i.\text{data}$.

Definition 5 (Collection Type). Two collections C_i and C_j have the same type (i.e., $\text{type}(C_i) = \text{type}(C_j)$) if any member e_i from C_i and any member e_j from C_j have the same type: $\text{type}(e_i) = \text{type}(e_j)$.

Composite. A group of visual elements that is neither a glyph nor a collection is considered a composite. Figure 6 shows two example composites: (a) a scatter plot matrix (SPLOM), where each member (scatter plot) is a collection of circle marks, and (b) a dashboard consisting of two collections visualizing different fields of the same dataset: a scatter plot on top and small multiples of waffle charts. The waffle chart small multiples is also a nested collection. Note that the composite in either Fig. 6a or Fig. 6b, is not a nested collection because condition 3 in Definition 4 is not satisfied. Members in a composite can be juxtaposed (e.g., Fig. 6b) or superimposed (e.g., Fig. 7a, where a polyline and a collection of rectangles share the same x- and y-axes) [19].

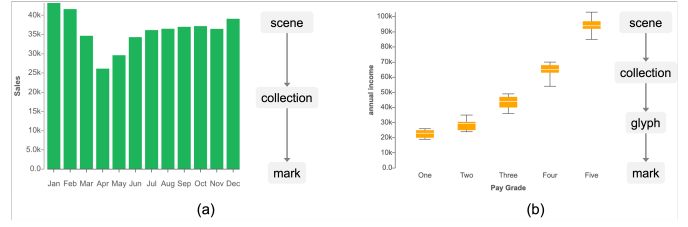


Fig. 5: (a) A collection of rectangles forms the main chart content of a bar chart; (b) a collection of glyphs forms that of a box plot.

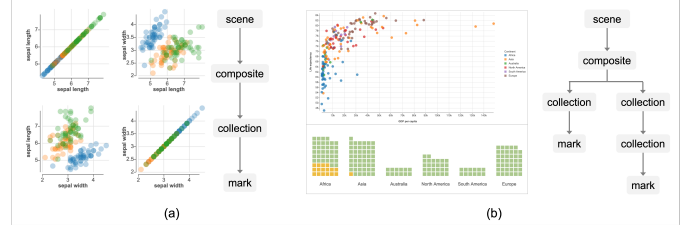


Fig. 6: (a) SPLOM is a composite of collections in the form of scatter plots, (b) a dashboard is a composite of multiple collections.

It is important to distinguish between glyph, collection, and composite because (1) no generative operation in MSC (Sec. 4.2) applies to all types of groups (e.g., we can classify a collection but not a glyph), and (2) knowing the relationships between different types of groups allows us to more precisely describe their creation and dissection (e.g., when a composite consists of multiple collections, we can selectively manipulate a subset of the collections).

3.4 Auxiliary Visual Elements

Auxiliary elements in a scene provide context and additional information for understanding the visualization.

Reference Element. A visualization scene can contain reference elements like axes, legends, grid lines and labels, whose primary purpose is to facilitate the reading and interpretation of the visual appearances of marks, glyphs and collections.

Annotation. Text, shapes, cues [21], and images [36] are often used as annotations to emphasize or explain specific data points, regions, or trends in a scene. Annotations provide additional context, offer clarity, and draw attention to specific parts of a visualization.

3.5 Visual Encodings and Scales

Visual encoding is a central component in data visualization, specifying the mapping of attribute values to visual channel values. In MSC, a visual encoding definition enc consists of a visual element e , a visual channel ch , a data attribute α , and a scale λ :

$$enc := \langle e, ch, \alpha, \lambda \rangle$$

The visual element in the encoding definition can be a mark, a vertex or segment, or a group. For example, in Fig. 1a the width (visual channel) of each rectangle mark (visual element) encodes pct (attribute); in Fig. 5b, the y coordinate (visual channel) of each top whisker segment (visual element) encodes the 75th percentile of annual income (attribute); in Fig. 7b, the x and y positions (visual channel) of each collection (visual element) encode the approximate geographic locations (attribute) of each state.

For each visual encoding, MSC requires the definition of a *scale*—a function specifying how the domain (data values) map to the range (visual channel values). For example, we have a scale in Fig. 1a that maps pct (attribute) values to the width (visual channel) of each rectangle mark: $\lambda(\text{pct value}) = \text{width value}$

Scales can be *shared* or *synced*. Figure 7a shows an example involving *shared* scales: the visualization scene is a composite consisting of a polyline and a rectangle collection; the y-position encoding of the

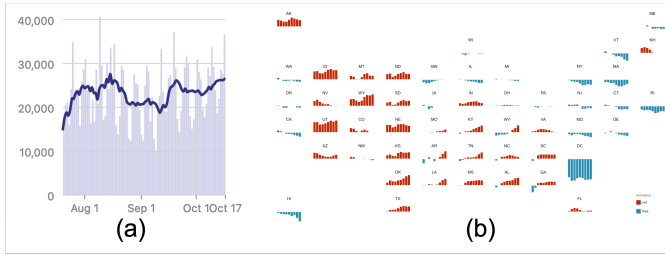


Fig. 7: (a) a line graph showing moving averages of COVID cases and a bar chart showing the raw case counts, (b) small multiples of bar charts showing how each state's partisan lean has shifted over the years, where each bar chart is positioned based on its approximate geographic location of the corresponding state.

vertices in the polyline shares the same scale as the height encoding of the bars; the x-position encoding of the polyline vertices shares the same scale as the x-position encoding of the bars. Two encodings can share the same scale as long as the attributes have the same data type, sharing a scale does not require two encodings having the same attribute or channels. Figure 7b shows an example involving *synced* scales. The visualization scene consists of 50 rectangle collections. Within each collection, the x-positions of the rectangles encode the year attribute. The scales across the 50 collections have the same domain (values of the year attribute), but the ranges expressed in terms of x coordinates are different. When editing this visualization, it is desirable that adjusting the scale range of one collection automatically updates the other collections to match. To achieve this, we can synchronize the scales across all 50 collections, so that updating the range extent of one scale will automatically adjust the range extents of the other scales.

3.6 Algorithmic Layouts

The spatial arrangement of visual elements is often achieved by visual encodings involving the position channels (i.e., x and y coordinates). In addition to such encodings, algorithmic layouts are often used to specify the positions of visual elements. Examples of algorithmic layouts include but are not limited to: grid (e.g., waffle chart), stack (e.g., stacked bar chart), packing [32] (e.g., beeswarm plot), and spiral (e.g., spiral plot). These layouts can be used to position not only marks and glyphs, but also collections (e.g., a composite in the form of a dashboard can use a tiling layout to arrange its member collections; a nested collection such as a scatter plot matrix applies a grid layout to its member collections). Each of these algorithmic layouts has a set of associated *parameters*: for example, a grid layout can specify the number of elements in a row/column and the direction of flow (e.g., top to bottom or left to right).

3.7 Relational Constraints

MSC includes relational constraints to specify inter-element relationships. The alignment constraint, for example, makes sure the anchors of elements are arranged in a straight line. In Fig. 1a, the light blue bars are aligned to the right, representing the divergence of opinions. To achieve this spatial arrangement, we can either align all the light blue bars to the right or all the light orange bars to the left. Another type of constraint is the affixation constraint, which specifies the position of one element relative to an anchor element. In Fig. 1a, we can describe the positions of the white text labels as a result of affixing the text marks at the center of the rectangles.

In addition, it is common to find constraints that specify the ordering of visual elements in a data visualization scene. For example, in a waffle chart, the ordering of rectangle marks in the collection together with the grid layout parameters determines the eventual visual appearance; in a connected scatter plot, the order of vertices on a polyline is determined by a temporal attribute and determines how the vertices are connected by segments. Ordering can also be specified along the z-dimension, which determines the order in which overlapping visual elements are drawn.

3.8 View Configuration

The view configuration component provides a vantage point of a data visualization scene with configurable attributes like focus point (center of the view port), field of view (extent visible through the view port), zoom level, and rotation. Visualizations such as MatrixWave [53] and GeneaQuilts [5], for example, apply rotation to the scenes.

3.9 Alternative Scene Descriptions

Three of the semantic components described above can determine the spatial positions of visual elements: visual encodings, algorithmic layouts, and relational constraints. It is possible to describe the same data visualization scene in multiple ways, each involving a different set of components. For example, we can describe the vertical bar chart in Fig. 5a as a collection of rectangles with a grid layout (one row, 12 columns) where the rectangles are sorted by the month attribute, or a collection of rectangles whose x positions are bound to the month attribute. Both descriptions are valid, but we should not include both the grid layout and the x encoding in the same description to describe the spatial arrangement of the bars along the horizontal direction.

4 PROCEDURE: GENERATIVE & MODIFICATIVE OPERATIONS

Given the definitions of semantic components in Sec. 3, MSC includes two kinds of operations on these components: 1) *generative* operations that create new visual elements (i.e., marks, glyphs, collections, composites) or remove existing ones based on data items and/or links, and 2) *modificative* operations that modify existing visual elements' properties or change the relationships between them.

Sec. 4.1 and Sec. 4.2 discuss generative operations: each operation consists of a descriptive name, an input visual element, an output element, and parameters. Most generative operations in MSC require all these parts to be specified. The input element and parameters may be optional for a few operations in Sec. 4.1. We use the following form to represent a generative operation, similar notations have been used by generative grammars in linguistics [12] and shape modeling [31].

`name(parameters) : input element \leadsto output element`

Sec. 4.3, Sec. 4.4, and Sec. 4.5 discuss modificative operations, where each operation is expressed either as *value assignment* in the form of `element.property = value`, where a visual element's property is set to a constant, or as a *mathematical constraint* in the form of `element1.property = f(element2.property, element3.property, ...)`, where a function *f* defines the relationships between the property values of multiple elements, usually expressed using arithmetic operators.

4.1 Generative Operations: Initialize Data and Elements

The first steps of creating a data visualization involve initializing a scene and importing data. In the operations below, `:=` denotes "defined as", `+` denotes "one or more", and `|` denotes "or":

`createScene() : \leadsto scn`

`importData(csvFile) : \leadsto table := items and attributes`

`importData(jsonFile) : \leadsto network := items, attributes and links`

We can then create primitive marks by specifying their types and properties, a scene is required as the input object:

`createMark(type, props) : scn \leadsto mark := rect|circle|line|path|txt|...`

One or more marks can be grouped to create a glyph:

`createGlyph() : mark+ \leadsto glyph`

4.2 Generative Operations: Join Visual Elements with Data

With primitive marks and data tables, we can now derive additional types of marks and create collections through graphics-data join operations. MSC specifies six such operations: repeat, divide, densify, classify, repopulate, and stratify.

Repeat. The repeat operation takes a mark, a glyph, or a collection as input and creates a collection of marks, glyphs, or collections (Fig. 8). This operation accepts a dataset and an ordinal or nominal attribute name in the dataset as its parameters. For example, in Fig. 2b, we repeat a rectangle by the age attribute. Since there are four unique age values in the data (below 30, 30 - 50, 50 - 70, above 70), we get a collection of

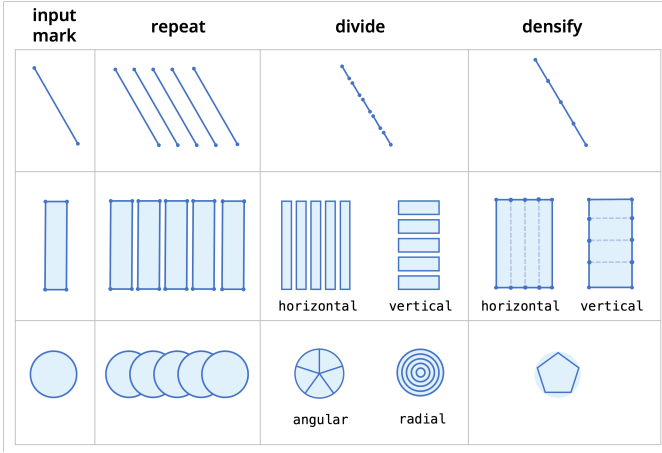


Fig. 8: Three operations for graphics-data join: repeat, divide, and densify. Here each input mark is joined with five unique attribute values.

four rectangles (Fig. 2b). The data scope of the first rectangle consists of all the rows with “below 30” as the age value, and so on.

Formally, when the dataset is a table, the output of a repeat operation is a collection $C = \{e_1, e_2, \dots, e_n\}$, where each member visual element e_i has the same type as the input element. The number of elements in C is equal to the number of unique values in the specified parameter attribute α . For each unique value v_i in α , there exists a visual element e_i in C such that $e_i.data[\alpha] = v_i$.

`repeat(table, attr) : mark \rightsquigarrow mark collection`
`repeat(table, attr) : glyph \rightsquigarrow glyph collection`
`repeat(table, attr) : collection \rightsquigarrow nested collection`

When the dataset is a network, let nd denote a visual element that represents an item, and lk denote a mark (usually a line, a path, an arc, or a band) that represents a link:

`repeat(network, attr) : [nd, lk] \rightsquigarrow [nd collection, lk collection]`

Divide. The divide operation splits a mark into a collection of smaller marks. This operation accepts a data table, and an ordinal or nominal attribute name in the table as its parameters. An additional orientation parameter is needed when the input mark type is rectangle or circle. The mark type in the output collection depends on the input mark type and the orientation.

`divide(table, attr, orientation) : rect \rightsquigarrow rect collection`
`divide(table, attr, orientation) : circle \rightsquigarrow pie collection | ring collection`
`divide(table, attr) : line \rightsquigarrow line collection`
`divide(table, attr) : pie \rightsquigarrow arc collection`
`divide(table, attr) : ring \rightsquigarrow arc collection`

For example, applying divide on a circle along the *angular* orientation produces a collection of pies, and along the *radial* orientation produces a collection of rings (Fig. 8). The graphics-data joining mechanism is similar to that of the repeat operation: tuples sharing the same attribute value are assigned as the data scope of each mark in the output collection.

Densify. In a line chart, area chart, or radar chart, we need to assign data scopes to the vertices of the geometric marks. MSC introduces the densify operation to facilitate the generation of vertices and the computation of data scopes. The densify operation transforms a rectangle into an area mark, a circle into a polygon, and a line into a polyline (Fig. 8) by adding vertices along the border of the input mark, assigning a data scope to each vertex, and replacing curve segments with line segments. This operation accepts a data table and an ordinal or nominal attribute name in the table as its parameters. An additional orientation parameter is needed when the input mark is a rectangle: vertices are added along parallel borders based on the specified orientation.

`densify(table, attr) : line \rightsquigarrow polyline`
`densify(table, attr) : circle \rightsquigarrow polygon`

`densify(table, attr, orientation) : rect \rightsquigarrow area`

The graphics-data joining mechanism is similar to that of the repeat operation and the divide operation. Each vertex represents a unique value in the attribute.

Classify. The classify operation groups visual elements inside a collection into multiple collections. This operation accepts an ordinal or nominal attribute name as its parameter.

`classify(attr) : collection \rightsquigarrow nested collection`

In Fig. 9, we classify a collection of eight rectangle marks by the region attribute, resulting in a nested collection consisting of four member collections, each of which consists of two rectangle marks.

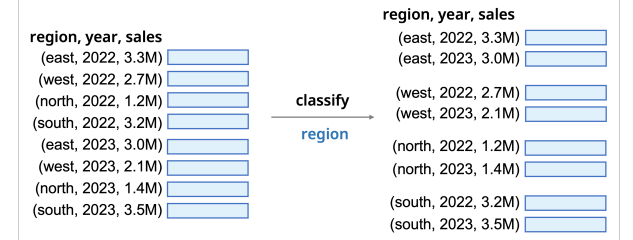


Fig. 9: Classifying a collection resulting in a nested collection. The text on the left of each mark denotes its data scope.

Formally, given an input collection $C = \{e_1, e_2, \dots, e_n\}$, classifying its members by an ordinal or nominal attribute α results in a nested collection, where the members of the input collection C are now collections: $C = \{c_1, c_2, \dots, c_k\}$. For each unique value v_i in the specified parameter attribute α , there exists a visual element c_i in C such that $c_i.data[\alpha] = v_i$. Each c_i contains some of the original members (i.e., e_1, e_2, \dots, e_n) such that for every member e_j in c_i , $c_i.data[\alpha] = e_j.data[\alpha]$.

Repopulate. A scene can be treated as a template, where the visualization design is applied to a new dataset. The repopulate operation enables such functionalities by re-generating visual elements inside an existing collection based on a new attribute in a new dataset. This operation accepts pairs of ordinal/nominal attribute names (one from the new dataset and the other from the existing dataset) as its parameters.

`repopulate({(attrnew, attrcurrent)}, tablenew) : collection \rightsquigarrow collection'`

Figure 10 shows how to change the underlying dataset of a grouped chart from a table about regional sales to a new table about countries' populations using the repopulate operation. The text on the left of each mark denotes its data scope.

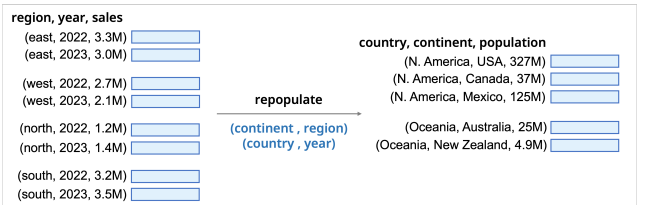


Fig. 10: Repopulating a nested collection where the number of member marks in the result changes based on the new dataset.

Formally, given an input collection $C = \{e_1, e_2, \dots, e_n\}$, after applying the repopulate operation by an attribute α from dataset T , we have a new collection $C' = \{e'_1, e'_2, \dots, e'_m\}$. Every element e'_j in C' has the same type as any member e_i in the input collection C , and the number of elements in C' is equal to the number of unique values in α . For each unique value v_i in α , there exists a visual element e'_i in C' such that $e'_i.data[\alpha] = v_i$.

Stratify. The stratify operation takes a mark as input and generates a collection of marks to be used as a layered representation of hierarchical data. This operation accepts a tree dataset (a special type of network dataset) and an ordinal or nominal attribute in the dataset as its parameters. The layering of generated marks is determined by the links connecting the items in the tree dataset. When the input mark is a

rectangle, an orientation parameter is needed to specify the direction of layering. Figure 11 shows that stratifying a rectangle and a circle results in an icicle plot [22] and a sunburst chart [45] respectively.

```
stratify(tree, attr, orientation) : rect  $\rightsquigarrow$  rect collection
stratify(tree, attr) : circle  $\rightsquigarrow$  arc collection
```

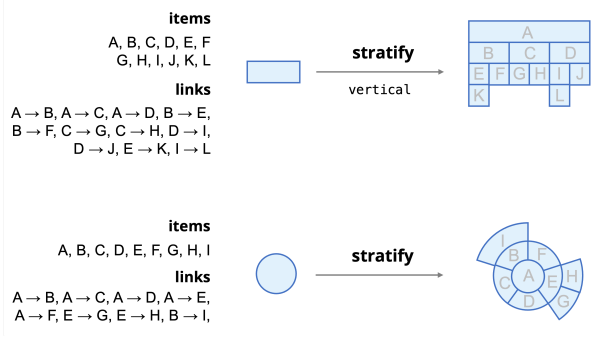


Fig. 11: Applying the stratify operation to a rectangle mark or a circle mark for a tree dataset.

4.3 Modificative Operations: Specify Channel Values

Apply/Remove Visual Encoding. Mapping attribute values to visual channel properties is a fundamental task in data visualization. For instance, the width of each rectangle mark in Fig. 2f encodes the pct value in the mark's data scope. Specifying such an operation in MSC requires the concept of *peers*. The peers of a visual element e , denoted by P_e , is a set of all the visual elements generated together with e through the repeat, divide, densify, classify, repopulate, and stratify operations in Sec. 4.2. For instance, all the four rectangle marks generated by repeat in Fig. 2b are peers of each other; all the sixteen rectangle marks generated by divide in Fig. 2c are peers of each other; all the vertices generated by a densify operation are peers of each other. Binding a data attribute α with a channel of a visual element e based on a scale λ is expressed as:

$$\forall e_i \in P_e, e_i.\text{channel} = \lambda(e_i.\text{data}[\alpha])$$

In MSC, the element e can be a vertex, a segment, a mark, a glyph, or a collection. This operation applies to both the initial mapping declaration as well as the update of channel values when the underlying data values for an existing mapping change.

Customize Scale. A scale directly determines how data attribute values are transformed into visual channel values. Assuming the data attribute values are fixed, we can specify the desired channel values by customizing the type (e.g., linear, power, log), domain (e.g., start from zero), and range (e.g., range extent for size, scheme for color) of scales.

Set Channel Values Directly. In the absence of any data attribute bound to a visual element's channel, the value of that channel can be freely changed. It is a common practice to modify the channel values for all the marks in a data visualization. For instance, the height of each rectangle mark in Figure 2(f) can be set to a constant k . Such an operation in MSC is expressed as:

$$\forall e_i \in P_e, e_i.\text{channel} = k$$

In certain cases, we may want to set the channel value for a specific visual element (e.g., for annotation purposes, to translate a visual element in design exploration) without affecting its peers, such an operation is expressed as: $e.\text{channel} = k$

4.4 Modificative Operations: Apply Layouts & Constraints

Apply Layout. In MSC, layouts are properties of visual element groups, and can be directly changed. A layout type and a set of associated parameters are required for a group G :

$$G.\text{layout} = \{\text{type}, \text{parameters}\}$$

Similar to the operation to set channel values, we can set the layout of multiple groups that are peers to each other:

$$\forall G_i \in P_G, G_i.\text{layout} = \{\text{type}, \text{parameters}\}$$

Update Layout Parameter. A parameter θ of an existing layout can be dynamically set to a new value v . For example, we can change the orientation of a stack layout from horizontal to vertical, or change the number of rows in a grid layout:

$$G.\text{layout}.\theta = v$$

Apply Relational Constraint. This operation applies to both initializing a new relational constraint as well as updating the channel values to maintain an existing constraint. Relational constraints in MSC are expressed using arithmetic operators. For example, to align the left sides of two marks' bounding boxes:

$$m_i.\text{bbox}.\text{left} = m_j.\text{bbox}.\text{left}$$

To affix the bottom left corner of a mark's bounding box with reference to the top right corner of another mark's bounding box:

$$m_i.\text{bbox}.\text{left} = m_j.\text{bbox}.\text{right} + dx$$

$$m_i.\text{bbox}.\text{bottom} = m_j.\text{bbox}.\text{top} + dy$$

In data visualization, relational constraints are usually applied to multiple pairs of visual elements. For instance, to attach a text label to the center of every rectangle mark in Fig. 1a:

$$\forall \text{rect}_i \in P_{\text{rect}}, \forall \text{text}_i \in P_{\text{text}},$$

$$\text{rect}_i.x = \text{text}_i.x \text{ and } \text{rect}_i.y = \text{text}_i.y,$$

$$\text{where } \text{rect}_i.\text{data} = \text{text}_i.\text{data}$$

4.5 Modificative Operations: Configure View

Finally, view configuration of a scene can be expressed by setting its properties, for example:

$$\text{scene.view.rotation} = 90^\circ$$

$$\text{scene.view.zoom} * = 1.2$$

4.6 Chaining Operations

A chart construction or manipulation task often requires chaining multiple operations described above. Here we provide two examples where the operations can be combined to achieve higher-level tasks.

Chaining Generative Operations. Concatenating these graphics-data join operations will produce nested structures. For instance, in Fig. 1, we perform repeat first followed by divide to obtain a stacked bar structure; applying the densify operation on a rectangle followed by the repeat operation yields a structure for small multiples of area charts.

Chaining Modificative Operations. Modifying the design of a chart may involve more than one operation. For instance, we may want to change the vertical bar chart in Fig. 12a into a horizontal bar chart in Fig. 12e, which can be accomplished through the following sequence of operations: Set Bar Height to Constant (Fig. 12b), Bind Attribute to Width (Fig. 12c), Update Grid Layout Parameter (Fig. 12d), and Customize Width Scale (Fig. 12e).

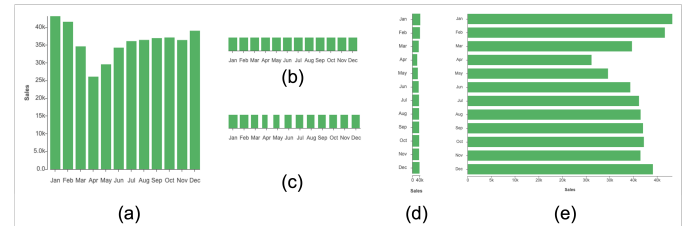


Fig. 12: Chaining modificative operations to change the orientation of a bar chart.

5 IMPLEMENTATION: MASCOT.JS

We implemented a JavaScript library, Mascot.js, based on the theoretical descriptions on the components (Sec. 3) and operations (Sec. 4) in MSC. A catalog of visualization examples created using Mascot.js is available at <https://mascot-vis.github.io>, which showcases more than 60 unique data visualization designs, covering commonly used chart types as well as more esoteric designs like cycle plot [39] and treemap bar chart.

API Design. Mascot.js provides global functions to import datasets, initialize scene objects, and create layouts. The generative operations are invoked as scene methods, operating within the context of that scene. Modificative operations applied to a single element are invoked through the element’s setter method, whereas modificative operations applied to multiple elements or an element’s peers are invoked through the appropriate methods in the scene object. The details on the API design are available in our short paper on a prior implementation of MSC called Atlas.js [24].

Smart Defaults. Mascot.js chooses smart default values if certain properties or components are not specified. For example, when creating new marks, default values for various channels are applied depending on the mark type. When a divide is performed, Mascot.js will apply a default layout to the resulting collection depending on the mark type. For instance, dividing a rectangle will also apply a stack layout to the generated smaller rectangles.

Creating Auxiliary Elements. Mascot.js provides convenience methods to create axes and legends. Users only need to specify the channel, attribute, and preferred properties such as position and visibility.

Rendering. Unlike libraries such as D3 [8] where visualization objects are represented as Document Object Model (DOM) nodes, Mascot.js represents visualization scenes as JavaScript objects and decouples scene abstractions from rendering. All the objects in a scene can be printed to the console and inspected without rendering. To display these objects, we need to create renderers. Mascot.js supports two types of renderers: SVG and WebGL.

Scene Serialization and Reconstruction. Scenes created using Mascot.js are JavaScript objects and can be serialized as JSON files. Mascot.js provides the functionality to ingest such JSON files and reconstruct the scene objects so that they can continue to be modified using MSC operations. The JSON files can also serve as templates: new visualizations can be created based on new datasets by applying a series of repopulate and apply visual encoding operations on the templates.

6 APPLICATIONS

In this section, we examine the potential of MSC as a foundation for three applications through case studies: interactive authoring of static representations, deconstructing and reusing existing charts, and augmenting static visualizations with animated behaviors.

6.1 Interactive Chart Authoring

Multiple research projects seek to design and build interactive authoring tools for users to construct static data visualizations in a WYSIWYG (What You See is What You Get) interface [20, 27, 37, 38, 40, 48]. To demonstrate the viability of MSC in powering real-world authoring applications, we implemented a new version of Data Illustrator [27] based on Mascot.js, which is available at <https://data-illustrateur.github.io/>. The benefits of using MSC as the core computational representation are twofold. First, thanks to a more complete theoretical formulation offered by MSC on data visualization components and operations, the new version of Data Illustrator is more expressive, supporting a range of data visualization designs that were not possible in the original version, including but not limited to rose chart, area chart, icicle plot, sunburst chart, and network diagrams. Second, MSC abstracts away the low-level complexities in scene manipulation and reduces development time and effort. The original Data Illustrator was built on Paper.js [23], which provides general scene graph abstraction support without visualization-specific components and operations. Compared to this original implementation involving 22632 lines of code in JavaScript, the new Data Illustrator implementation based on Mascot.js has only 5519 lines of code—more than four times reduction.

6.2 Chart Deconstruction and Reuse

Since constructing visualizations from scratch can still require significant knowledge and skills in using a particular authoring tool, researchers have also investigated methods and systems that automatically analyze existing visualizations and transform them into reusable templates [9, 11, 13, 17, 18, 33, 34].

We investigated the potential of MSC as a computational representation for visualization reuse in the Mystique system [9]. Mystique takes existing rectangle-based visualizations in the SVG format as input, and extracts scene components including reference elements (axis & legend), groups (collection & glyph), algorithmic layouts (grid, stack, packing), and relational constraints (alignment & affixation). The extracted components can then be manipulated and infused with a new dataset to produce a new scene through operations such as repopulate, apply visual encoding, and apply layout & constraint. Mystique provides a wizard interface to guide users in specifying the mappings between the components and the new dataset, and produce immediate visual feedback by updating the input example with the incrementally specified mappings. The benefits of using MSC as the computational representation for visualization deconstruction and reuse include the following: 1) we can directly establish correspondence between SVG elements and scene components, without having to go through an intermediate formal language, 2) the direct representation of scenes allows us to dive deep into complex layout structures such as nesting and small multiples, and 3) the intermediate scene results generated in the reuse process do not have to conform to the semantics and syntax of a formal language, allowing users focus on one mapping at a time.

6.3 Animating Static Visualizations

Besides authoring and reusing static visualizations, there is growing interest in research works to augment static visualizations with interactive and animated behaviors [14, 15, 44, 47, 50, 54]. To enable such augmentation, there exists two approaches to ingest the static visualizations. The first approach requires that the static visualizations must be created by a particular tool. For example, Data Animator [47] adds animated behaviors to static visualizations created using Data Illustrator [27], and Lyra 2 [54] turns static visualizations composed using Vega-Lite [42] into interactive ones. Such assumptions ensure that the semantic components are readily usable for behavior specification. Alternatively, we can broaden the scope to support input static visualizations in a specific format generated by multiple tools, but the visualizations need to be annotated with semantic component labels. For instance, CAST [14] is an animation authoring tool and works with SVG charts, provided that these charts are annotated with labels in a format called dSVG (data-enriched SVG) [15]. The dSVG format used by CAST can be considered an attempt to formulate a computational representation, however, it is not well documented and is a one-time solution for this specific task of animation authoring.

We are interested in understanding how well a more generic computational representation like MSC supports animation behavior specification. In this case study, we focus on two existing systems: Data Animator and CAST. As mentioned earlier, Data Animator requires static visualizations to be created using Data Illustrator, and MSC’s component formulation is a more complete account and encompasses the component model of Data Illustrator, it is thus trivial to convert visualizations expressed using MSC into the Data Illustrator format. CAST, on the other hand, uses dSVG as the underlying computational representation. Unlike MSC, dSVG does not explicitly describe the hierarchical relationships between scene elements, and adds three properties “id”, “class”, and “datum” for each mark. Based on the grouping and data scope components in MSC, we were able to develop scripts that convert data visualization scenes created using Mascot.js into the dSVG format, and verified that the resulting scenes in the dSVG format worked well with the CAST system. The supplemental materials contain the conversion scripts as well as video demos demonstrating the specification of animation behaviors. This exercise validates MSC as a computational representation for augmenting static visualizations, demonstrating its compatibility with current animation authoring tools and its potential as a foundation to develop similar systems in the future.

7 COMPARISON WITH RELATED FRAMEWORKS

As discussed in the introduction, the goal of proposing MSC as a new computational representation for data visualization scenes is to support various applications for constructing, decomposing, and augmenting visualizations. While Mascot.js can be considered yet another data

Table 1: A comparison of the primitives (i.e., components and operations) defined in various data visualization frameworks. A ✓ means the primitive is included in the framework, but the paper does not formally name it. If a primitive is named, we include the name used in the paper in the corresponding cell. A ✗ means the primitive is not included in the paper.

	MSC	Protovis [7]	D3 [8]	Vega [3]	Data Illustrator [27]	Chartculator [38]	Bluefish [35]	Chartreuse [13]	AutoTimeline [11]	InfoMotion [50]
Components	mark	mark	mark	mark	mark	mark	element	visual element	mark	visual element
	data scope	✗	data	✗	data scope	✓	✗	✗	✗	✗
	glyph	✗	✗	✗	group	glyph	group	unit	✗	repeating unit
	collection	panel	✗	✗	collection	✓	group	✗	layout	group
	composite	panel	✗	✗	✗	✗	✗	✗	✗	✗
	reference element	rule	axis, legend	axis, legend	axis, legend	axis, legend	✗	✗	✗	✗
	annotation	✗	✗	✗	✗	✓	connectedness, common region	embellishment	annotation	connectors
	visual encoding	✓	✗	✓	✓	visual encoding	✓	visual element	visual encoding	✗
	scale	scale	scale	scale	✓	scale	✗	✗	scale	✗
	algorithmic layouts	layouts	layouts	layouts	layouts	constraint-based layout	✗	unit layout	representation	unit layout
Operations	relational constraints	anchor	✗	✗	✓	✓	perceptual relation	unit layout	✗	unit layout
	view configuration	✗	✗	✓	✗	✗	✗	✗	✗	✗
	repeat	✗	join	✗	✓	✓	✓	✗	✗	✗
	divide	✗	✗	✗	✓	✗	✗	✗	✗	✗
	densify	✗	✗	✗	✗	✗	✗	✗	✗	✗
	classify	✗	✗	✗	✗	✗	✗	✗	✗	✗
	repopulate	✗	✗	✓	✗	✓	✗	repeat	✓	✗
	stratify	✗	✗	✓	✗	✗	✗	✗	✗	✗
	apply visual encoding	✓	✓	✓	✓	✓	✓	morph, recolor	✓	✗
	customize scale	✓	✓	✓	✓	✓	✗	✗	✗	✗
	set channel values	✓	✓	✓	✓	✓	✓	fix	✗	✗
	apply layout	✓	✓	✓	✓	✓	✓	partition	✓	✗
	update layout parameters	✓	✓	✓	✓	✓	✓	✗	✗	✗
	apply relational constraint	✓	✗	✗	✓	✓	✓	move	✗	✗
	configure view	✗	✗	✓	✗	✗	✗	✗	✗	✗

visualization library, its usability and learnability as a domain-specific language are not our primary concern. Instead, we are interested in understanding the strengths and weaknesses of MSC as a theoretical framework in enabling various applications. To evaluate MSC, we first present an overall comparison with nine visualization frameworks in terms of component and operation abstractions. We then focus on more in-depth comparisons with three classes of related work: DOM manipulation libraries (D3), declarative languages (e.g., Vega/Vega-Lite), and algebraic frameworks (e.g., VizQL).

7.1 Component and Operation Abstractions

Many visualization libraries and tools have proposed scene abstractions for various applications. Protovis [7] argues that designers should be able to think in terms of graphical marks, not abstract specifications to improve accessibility, and MSC shares the same stance. The components and operations in MSC are directly informed by recent visualization authoring frameworks and tools [41]. Specifically, the graphics-centric approach in MSC is consistent with the lazy data binding approach in Data Illustrator [27], and a few operations in MSC are derived from the repeat and partition operators in Data Illustrator [27]; Chartculator’s constraint-based layout approach [38] inspired the modificative operations in MSC.

To assess the completeness of abstraction, we compare MSC with nine visualization frameworks (Tab. 1), focusing on visual representations instead of interaction or animation. Among these, six were developed primarily for visualization creation: Protovis [7], D3 [8], Vega [3], Data Illustrator [27], Chartculator [38], and Bluefish [35]; two were developed primarily for infographics reuse [11, 13]; and one was developed for augmenting visualization with animation [50]. In general, MSC offers a more complete set of components and operations (Tab. 1). It is important to note that completeness is not equivalent to expressiveness: a missing component or operation in a framework does not mean that the corresponding visualization feature cannot be realized using the framework. For example, D3 provides very few abstractions but is highly expressive due to its reliance on the DOM as the scene representation. In addition, due to the different philosophies and approaches adopted by these frameworks, a direct one-to-one comparison is not always possible. For example, Vega [3] does not provide some generative operations because it is essentially a declarative specification language, not a scene representation. We elaborate on the differences

between MSC and D3 [8] in Sec. 7.2 and MSC and Vega in Sec. 7.3.

7.2 Data-Driven Documents (D3)

The main goal of D3 [8] is to support data-driven manipulation of documents represented using the DOM (document object model) standard. D3 provides very limited scene abstraction for data visualization, offering only five types of components (mark, data, axis/legend, scale, and layout). Everything else needs to be represented and manipulated as SVG elements through low-level methods. Such an approach enables an unparalleled level of expressiveness compared to other visualization frameworks. However, neither the low-level JavaScript code nor the resultant SVG scene graph serve as effective computational representations for applications beyond scene assembly. For instance, D3 provides no encapsulation for components like visual encodings or relational constraints: elements’ visual properties need to be computed using low-level code, and any subsequent manipulation of the SVG scene graph will not automatically update the element properties to preserve the encodings or constraints. The lack of higher-level abstraction makes it challenging to build graphical user interfaces on top of D3 to support interactive authoring from scratch; to date, no such tools have been created. Similarly, without high-level encapsulation of components and operations, efforts to deconstruct and reuse D3 visualizations mostly focus on visual styles [17, 18], and cannot be easily extended to structural features like layouts [9].

7.3 Declarative Specifications

Compared to D3, visualization grammars such as ggplot2 [51] and Vega [42, 43] provide higher-level and more systematic abstractions. In these approaches, the abstractions are in the form of declarative specifications that succinctly describe the data transformations and visual encodings in a visualization. Most of these declarative specifications are implementations of the Grammar of Graphics [52], which defines grammatical components as “rules for constructing graphs mathematically”. In contrast, MSC describes the semantic structures and operations from a graphics-centric perspective. For example, the Grammar of Graphics describes a pie chart as the result of mapping data categories to rectangles and performing a polar transformation [52]; under the MSC paradigm, we may describe the process as dividing a circle and assigning categories to each pie, which is more interpretable.

Unlike MSC and D3, declarative specifications are not representations of scenes. They are descriptions of what the scenes should be like—as Wilkinson put it, “A scene and its description are different” [52] (p.7). Specifications need to be compiled into scenes, which are typically represented and rendered using SVG. While grammars like Vega do have internal representations of scenes, they hide details such as the structural relationships between visual elements from users. Any modification to the scene must be done by changing the specification, and users have no means to select or modify scene components directly. We elaborate on the implications of this difference for two sample applications below:

Interactive Authoring. In Lyra [40], Vega serves as the underlying computational representation: user interaction in the GUI translates to changes in an underlying specification, which is dynamically compiled to a visualization scene (Fig. 13a). In comparison, MSC provides a more direct experience in scene manipulation where the intermediate specifications are eliminated (Fig. 13b). The MSC approach has several potential benefits: 1) during the authoring process, users can more freely manipulate the visualization scene, which may not conform to the semantics and syntax of a declarative specification language. Consequently, the authoring interface affords less restricted creative exploration and tinkering [28], 2) the explicit conceptualization of semantic components such as collection and constraints afford direct manipulation interaction designs, and 3) there is no need to re-compile the entire specification every time a minor scene modification happens, which can lead to a more responsive and lower-latency experience.

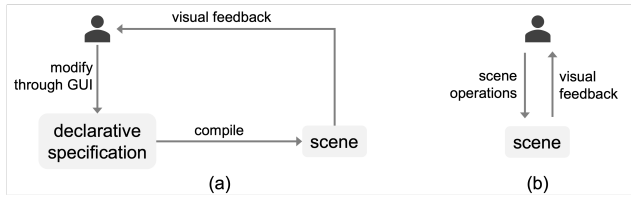


Fig. 13: Interactive authoring using (a) declarative specifications and (b) MSC as the underlying computational representations.

Visualization Deconstruction and Reuse. A typical processing pipeline for this application involves two stages: 1) reverse engineer an existing data visualization to understand its structure, 2) generate a new visualization of a new dataset based on this understanding. The input visualizations are usually in a raster image or vector graphics format. The output of the reverse engineering stage and the input of the generation stage can be expressed either as declarative specifications (e.g., [33]) or scene abstractions (e.g., [9, 11, 13]). In Mystique [9], we demonstrated how the reconstruction results for SVG charts with complex layout structures can be directly expressed using MSC, and the scene components can be dynamically updated through MSC operations to infuse with a new dataset (Fig. 14b). If declarative specifications were instead chosen as the computational representation, deconstruction and reuse would be more challenging: we would need to convert the deconstruction results into an intermediate specification (Fig. 14a), and dynamically generate partial specifications during the reuse process. Currently, no such solutions exist for charts with advanced layouts like small multiples and nested grouping.

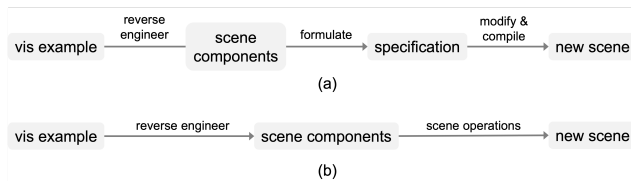


Fig. 14: Deconstructing and reusing existing visualizations using (a) declarative specifications and (b) MSC as the underlying computational representations.

7.4 Algebraic Frameworks

VizQL [16] is a formal language used by Tableau [2] to describe data visualizations. VizQL allows the specification of table configurations (i.e., rows and columns) as well as visual encodings within each pane. Specifically, table configurations are expressed using an algebra, where the operands are attributes (categorized into dimensions and measures) and the operators include cross, next, and concatenation [46].

The operations in MSC can be considered algebraic: the computation of data scopes when chaining the repeat, divide, and densify operations is equivalent to the nest operator in VizQL; the classify operation in MSC is similar to the group by operation in relational algebra; and the repeat operation for network data is equivalent to the concatenation of the create nodes and create connections in the Ploceus framework [25, 26]. However, in MSC operations, the primary operands are visual elements, and the attributes are secondary. The generative operations in MSC describe how marks and groups are derived from input visual elements; the modificative operations in MSC describe the algebraic relationships between visual channel properties.

The attribute-centric algebra in VizQL represents a top-down approach to visualization generation: one starts from table configurations, then proceeds to mark choice and visual encodings in each table pane. This approach enables easy and systematic specifications of complex nested designs like small multiples. In contrast, the graphics-centric algebra in MSC is bottom-up: we start from marks and glyphs, and proceed to higher-level groupings. Additional effort is needed to apply the appropriate layouts to achieve table-based designs. In VizQL, since visual elements are not treated as first-class citizens, describing visualizations that do not fit neatly into a row-and-column metaphor can be challenging. For example, specifying glyph-based plots (e.g., box plot, dumbbell chart) involves non-trivial workarounds with custom attribute derivation and calculations.

8 CONCLUSION AND FUTURE WORK

To support data visualization applications such as reverse engineering and authoring, a computational representation is needed to abstract the components and their manipulations in a visualization scene. Manipulable Semantic Components (MSC) provide this by specifying a unified object model for the components in a visualization scene and a set of operations to generate and modify the components.

In future work, we plan to support both design-time and run-time interactions in MSC. The current focus of MSC is on design-time manipulation (e.g., construction and editing of visual representations); however, letting users specify run-time interactive behaviors (e.g., brushing, filtering) tailored for different scenarios (e.g., exploratory analysis, scrolly-telling) is essential. Designing a unified architecture for both kinds of interactions remains a challenge. One potential solution is to model the relationships between the components as a dependency graph: whenever the value of a variable (e.g., a component’s channel) is changed, the architecture should automatically propagate the change and update the values of other dependent variables. To capture the dependencies, we may formulate the relationships between components as constraints [38, 49]: for example, the generative operations, layout operations, and encoding-related operations are one-way constraints, while operations that apply affixation and alignment can be described as multi-way constraints. Future research can focus on how to model and solve such constraint-based dependency graphs.

Another interesting future direction is to explore how to represent and learn design knowledge from visualizations expressed using MSC components. For example, Draco [29] formalize and extract design knowledge from declarative specifications, and we are interested in understanding if Draco can be easily extended to work on scene representations such as MSC, and if novel methods can be developed to model design knowledge.

ACKNOWLEDGMENTS

Zhicheng Liu and Chen Chen were supported in part by NSF grant IIS-2239130.

REFERENCES

- [1] Scalable Vector Graphics (SVG) 2. <https://www.w3.org/TR/SVG2/>. 1
- [2] Tableau: Business Intelligence and Analytics Software. 9
- [3] Vega: a visualization grammar, Apr. 2021. <https://github.com/vega/vega>. 8
- [4] Accurat. Brain drain, Jan. 2014. <https://www.behance.net/gallery/14187547/Brain-drain>. 3
- [5] A. Bezerianos, P. Dragicevic, J.-D. Fekete, J. Bae, and B. Watson. GeneaQuilts: A System for Exploring Large Genealogies. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1073–1081, Nov. 2010. doi: 10.1109/TVCG.2010.159 4
- [6] A. Bigelow, S. Drucker, D. Fisher, and M. Meyer. Iterating between Tools to Create and Edit Visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):481–490, Jan. 2017. doi: 10.1109/TVCG.2016.2598609 1
- [7] M. Bostock and J. Heer. Protovis: A graphical toolkit for visualization. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1121–1128, 2009. Publisher: IEEE. doi: 10.1109/TVCG.2009.174 8
- [8] M. Bostock, V. Ogievetsky, and J. Heer. D3: data-driven documents. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2301–2309, 2011. Publisher: IEEE. doi: 10.1109/TVCG.2011.185 1, 7, 8
- [9] C. Chen, B. Lee, Y. Wang, Y. Chang, and Z. Liu. Mystique: Deconstructing SVG Charts for Layout Reuse. *IEEE Transactions on Visualization and Computer Graphics*, 30(1):447–457, 2023. Publisher: IEEE. doi: 10.1109/TVCG.2023.3327354 7, 8, 9
- [10] C. Chen and Z. Liu. The State of the Art in Creating Visualization Corpora for Automated Chart Analysis. *Computer Graphics Forum*, 42(3):449–470, 2023. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14855>. doi: 10.1111/cgf.14855 1
- [11] Z. Chen, Y. Wang, Q. Wang, Y. Wang, and H. Qu. Towards Automated Infographic Design: Deep Learning-based Auto-Extraction of Extensible Timeline. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):917–926, Jan. 2020. doi: 10.1109/TVCG.2019.2934810 1, 7, 8, 9
- [12] N. Chomsky. *Syntactic Structures*. Walter de Gruyter, 2002. 4
- [13] W. Cui, J. Wang, H. Huang, Y. Wang, C.-Y. Lin, H. Zhang, and D. Zhang. A Mixed-Initiative Approach to Reusing Infographic Charts. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):173–183, Jan. 2022. Conference Name: IEEE Transactions on Visualization and Computer Graphics. doi: 10.1109/TVCG.2021.3114856 1, 7, 8, 9
- [14] T. Ge, B. Lee, and Y. Wang. CAST: Authoring Data-Driven Chart Animations. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21, pp. 1–15. Association for Computing Machinery, New York, NY, USA, May 2021. doi: 10.1145/3411764.3445452 7
- [15] T. Ge, Y. Zhao, B. Lee, D. Ren, B. Chen, and Y. Wang. Canis: A High-level Language for Data-Driven Chart Animations. *Computer Graphics Forum*, 39(3), 2020. doi: 10.1111/cgf.14005 1, 7
- [16] P. Hanrahan. VizQL: a language for query, analysis and visualization. In *Proceedings of the 2006 ACM SIGMOD international conference on Management of data*, pp. 721–721, 2006. doi: 10.1145/1142473.1142560 9
- [17] J. Harper and M. Agrawala. Deconstructing and restyling D3 visualizations. In *Proceedings of the 27th annual ACM Symposium on User Interface Software and Technology*, UIST '14, pp. 253–262. Association for Computing Machinery, New York, NY, USA, Oct. 2014. doi: 10.1145/2642918.2647411 7, 8
- [18] J. Harper and M. Agrawala. Converting basic D3 charts into reusable style templates. *IEEE Transactions on Visualization and Computer Graphics*, 24(3):1274–1286, 2017. Publisher: IEEE. doi: 10.1109/TVCG.2017.2659744 7, 8
- [19] W. Javed and N. Elmqvist. Exploring the design space of composite visualization. In *2012 IEEE Pacific Visualization Symposium*, pp. 1–8. IEEE, 2012. doi: 10.1109/PacificVis.2012.6183556 3
- [20] N. W. Kim, E. Schweickart, Z. Liu, M. Dontcheva, W. Li, J. Popovic, and H. Pfister. Data-driven guides: Supporting expressive design for information graphics. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):491–500, 2016. Publisher: IEEE. doi: 10.1109/TVCG.2016.2598620 7
- [21] H.-K. Kong, Z. Liu, and K. Karahalios. Internal and external visual cue preferences for visualizations in presentations. In *Computer Graphics Forum*, vol. 36, pp. 515–525. Wiley Online Library, 2017. Issue: 3. doi: 10.1111/cgf.13207 3
- [22] J. B. Kruskal and J. M. Landwehr. Icicle plots: Better displays for hierarchical clustering. *The American Statistician*, 37(2):162–168, 1983. Publisher: Taylor & Francis. 6
- [23] J. Lehn and J. Puckey. Paper.js: The Swiss Army Knife of Vector Graphics Scripting. 7
- [24] Z. Liu, C. Chen, F. Morales, and Y. Zhao. Atlas: Grammar-based Procedural Generation of Data Visualizations. In *2021 IEEE Visualization Conference (VIS)*, 2021. doi: 10.1109/VIS49827.2021.9623315 7
- [25] Z. Liu, S. B. Navathe, and J. T. Stasko. Network-based visual analysis of tabular data. In *2011 IEEE Conference on Visual Analytics Science and Technology (VAST)*, pp. 41–50, Oct. 2011. doi: 10.1109/VAST.2011.6102440 9
- [26] Z. Liu, S. B. Navathe, and J. T. Stasko. Ploceus: Modeling, visualizing, and analyzing tabular data as networks. *Information Visualization*, 13(1):59–89, Jan. 2014. Publisher: SAGE Publications. doi: 10.1177/1473871613488591 9
- [27] Z. Liu, J. Thompson, A. Wilson, M. Dontcheva, J. Delorey, S. Grigg, B. Kerr, and J. Stasko. Data Illustrator: Augmenting vector design tools with lazy data binding for expressive visualization authoring. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pp. 1–13, 2018. doi: 10.1145/3173574.3173697 7, 8
- [28] P. Louridas. Design as bricolage: anthropology meets design thinking. *Design Studies*, 20(6):517–535, 1999. Publisher: Elsevier. 9
- [29] D. Moritz, C. Wang, G. L. Nelson, H. Lin, A. M. Smith, B. Howe, and J. Heer. Formalizing visualization design knowledge as constraints: Actionable and extensible models in Draco. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):438–448, 2018. Publisher: IEEE. doi: 10.1109/TVCG.2018.2865240 9
- [30] T. Munzner. *Visualization analysis and design*. CRC press, 2014. 2
- [31] P. Müller, P. Wonka, S. Haegler, A. Ulmer, and L. Van Gool. Procedural modeling of buildings. In *ACM SIGGRAPH 2006 Papers*, SIGGRAPH '06, pp. 614–623. Association for Computing Machinery, Boston, Massachusetts, July 2006. doi: 10.1145/1179352.1141931 4
- [32] D. Park, S. M. Drucker, R. Fernandez, and N. Elmqvist. Atom: A Grammar for Unit Visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 24(12):3032–3043, Dec. 2018. Conference Name: IEEE Transactions on Visualization and Computer Graphics. doi: 10.1109/TVCG.2017.2785807 4
- [33] J. Poco and J. Heer. Reverse-Engineering Visualizations: Recovering Visual Encodings from Chart Images. *Computer Graphics Forum*, 36(3):353–363, 2017. doi: 10.1111/cgf.13193 7, 9
- [34] J. Poco, A. Mayhua, and J. Heer. Extracting and Retargeting Color Mappings from Bitmap Images of Visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):637–646, Jan. 2018. Conference Name: IEEE Transactions on Visualization and Computer Graphics. doi: 10.1109/TVCG.2017.2744320 7
- [35] J. Pollock, C. Mei, G. Huang, D. Jackson, and A. Satyanarayan. Bluefish: A Relational Framework for Graphic Representations, Nov. 2023. arXiv:2307.00146 [cs]. doi: 10.48550/arXiv.2307.00146 8
- [36] D. Ren, M. Brehmer, B. Lee, T. Höllerer, and E. K. Choe. ChartAccent: Annotation for data-driven storytelling. In *2017 IEEE Pacific Visualization Symposium (PacificVis)*, pp. 230–239. Ieee, 2017. doi: 10.1109/PACIFICVIS.2017.8031599 1, 3
- [37] D. Ren, T. Höllerer, and X. Yuan. iVisDesigner: Expressive interactive design of information visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2092–2101, 2014. Publisher: IEEE. doi: 10.1109/TVCG.2014.2346291 7
- [38] D. Ren, B. Lee, and M. Brehmer. Charticulator: Interactive construction of bespoke chart layouts. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):789–799, 2018. Publisher: IEEE. doi: 10.1109/TVCG.2018.2865158 7, 8, 9
- [39] N. B. Robbins. Introduction to Cycle Plots. *Visual Business Intelligence Newsletter*, 2008. 6
- [40] A. Satyanarayan and J. Heer. Lyra: An Interactive Visualization Design Environment. *Computer Graphics Forum*, 33(3):351–360, 2014. doi: 10.1111/cgf.12391 7, 9
- [41] A. Satyanarayan, B. Lee, D. Ren, J. Heer, J. Stasko, J. Thompson, M. Brehmer, and Z. Liu. Critical reflections on visualization authoring systems. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):461–471, 2019. Publisher: IEEE. doi: 10.1109/TVCG.2019.2934281 8
- [42] A. Satyanarayan, D. Moritz, K. Wongsuphasawat, and J. Heer. Vega-Lite: A grammar of interactive graphics. *IEEE Transactions on Visualization*

- and *Computer Graphics*, 23(1):341–350, 2016. Publisher: IEEE. doi: [10.1109/TVCG.2016.2599030](https://doi.org/10.1109/TVCG.2016.2599030) 1, 7, 8
- [43] A. Satyanarayan, R. Russell, J. Hoffswell, and J. Heer. Reactive Vega: A Streaming Dataflow Architecture for Declarative Interactive Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):659–668, Jan. 2016. Conference Name: IEEE Transactions on Visualization and Computer Graphics. doi: [10.1109/TVCG.2015.2467091](https://doi.org/10.1109/TVCG.2015.2467091) 8
 - [44] L. S. Snyder and J. Heer. DIVI: Dynamically Interactive Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 30(1):403–413, 2024. Publisher: IEEE. doi: [10.1109/TVCG.2023.3327172](https://doi.org/10.1109/TVCG.2023.3327172) 1, 7
 - [45] J. Stasko and E. Zhang. Focus+ context display and navigation techniques for enhancing radial, space-filling hierarchy visualizations. In *IEEE Symposium on Information Visualization 2000. INFOVIS 2000. Proceedings*, pp. 57–65. IEEE, 2000. doi: [10.1109/INFVIS.2000.885091](https://doi.org/10.1109/INFVIS.2000.885091) 6
 - [46] C. Stolte, D. Tang, and P. Hanrahan. Polaris: A system for query, analysis, and visualization of multidimensional relational databases. *IEEE Transactions on Visualization and Computer Graphics*, 8(1):52–65, 2002. Publisher: IEEE. doi: [10.1109/INFVIS.2000.885086](https://doi.org/10.1109/INFVIS.2000.885086) 9
 - [47] J. R. Thompson, Z. Liu, and J. Stasko. Data Animator: Authoring Expressive Animated Data Graphics. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI ’21, pp. 1–18. Association for Computing Machinery, New York, NY, USA, May 2021. doi: [10.1145/3411764.3445747](https://doi.org/10.1145/3411764.3445747) 7
 - [48] T. Tsandilas. StructGraphics: Flexible Visualization Design through Data-Agnostic and Reusable Graphical Structures. *IEEE Transactions on Visualization and Computer Graphics*, 2020. Publisher: IEEE. doi: [10.1109/TVCG.2020.3030476](https://doi.org/10.1109/TVCG.2020.3030476) 7
 - [49] B. Vander Zanden, R. Halterman, B. Myers, R. McDaniel, R. Miller, P. Szekely, D. Giuse, and D. Kosbie. Lessons learned about one-way, dataflow constraints in the Garnet and Amulet graphical toolkits. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 23:776–796, Nov. 2001. doi: [10.1145/506315.506318](https://doi.org/10.1145/506315.506318) 9
 - [50] Y. Wang, Y. Gao, R. Huang, W. Cui, H. Zhang, and D. Zhang. Animated Presentation of Static Infographics with InfoMotion. *Computer Graphics Forum*, 40(3):507–518, June 2021. doi: [10.1111/cgf.14325](https://doi.org/10.1111/cgf.14325) 1, 7, 8
 - [51] H. Wickham. *ggplot2: Elegant graphics for data analysis*. Springer, 2009. 1, 8
 - [52] L. Wilkinson. *The Grammar of Graphics*. Springer, New York, 2nd edition ed., July 2005. 8, 9
 - [53] J. Zhao, Z. Liu, M. Dontcheva, A. Hertzmann, and A. Wilson. MatrixWave: Visual Comparison of Event Sequence Data. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 2015. doi: [10.1145/2702123.2702419](https://doi.org/10.1145/2702123.2702419) 4
 - [54] J. Zong, D. Barnwal, R. Neogy, and A. Satyanarayan. Lyra 2: Designing interactive visualizations by demonstration. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):304–314, 2020. Publisher: IEEE. doi: [10.1109/TVCG.2020.3030367](https://doi.org/10.1109/TVCG.2020.3030367) 7