

# Database Theory in Action: Database Visualization

Eugene Wu   

Columbia University, New York, NY, USA

---

## Abstract

---

We draw a connection between data modeling and visualization, namely that a visualization specification defines a mapping from database constraints to visual representations of those constraints. We show data modeling explains many existing visualization design and introduce multi-table *database* visualization.

**2012 ACM Subject Classification** Human-centered computing → Visualization theory, concepts and paradigms; Information systems → Database design and models

**Keywords and phrases** Visualization Theory, Data Model, Database Visualization

**Digital Object Identifier** 10.4230/LIPIcs.ICDT.2025.35

**Funding** This material is based upon work supported by NSF 1845638, 1740305, 2008295, 2106197, 2103794, 2312991, Amazon, Google, Adobe, and CAIT. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the funders.

**Acknowledgements** Thanks to Wolfgang Gatterbauer and Eirik Bakke for early feedback on this work.

## 1 Introduction

Existing visualization theory, libraries, and graphical grammars assume that a single input table is processed to produce an output data visualization. Yet, database theory shows that one table does not model all data (without introducing redundancy). We go beyond the single-table data model towards a theory of multi-table *Database Visualization*. The theory models a visualization as a visual representation of database contents (the rows and values) *that preserves database constraints*. We present the main ideas through a series of examples and show how under this theory, some common visualization designs are a consequence of data modeling choices.

## 2 Background

This section reviews visualization theory, the relational data model, and single table grammars.

### 2.1 Visualization Theory and the Single-Table Assumption

Bertin first described the data-to-visual (or pixel) mapping, where tuples map to marks (e.g., points), and data attributes map to the mark’s visual channels [2]. For instance, Figure 1(a) visualizes a table  $T(\text{id}, a, b)$  as a scatter plot  $V_T$  where  $T.a$  and  $T.b$  are mapped to the x and y axes (ignore the red edges). Wilkinson’s Grammar of Graphics [8] extended this mapping with scales, mark types, facets, and more to develop a “graphical grammar” that defines a design space of data visualizations, which forms the foundation of most data visualization systems including nViZn [7], ggplot2 [6], Tableau [5], and vega-lite [4].

However, this theory is predicated on a single-table data model, which does not model most data such as hierarchical JSON, RDF graphs, and relational databases. Thus, the user is expected to “prepare” their data (extract, transform, filter, join) into a single table before visualization. This is a major burden, cannot guarantee that the visualization faithfully



© Eugene Wu;

licensed under Creative Commons License CC-BY 4.0


28th International Conference on Database Theory (ICDT 2025).

Editors: Sudeepa Roy and Ahmet Kara; Article No. 35; pp. 35:1–35:6

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 1** In the node link visualization, (a) the links  $V_{E'}$  only appear to connect the points  $V_T$ . (b) The points and links become inconsistent if  $V_T$  changes (e.g., jitter).

reflects the data, and makes it impossible for existing formalisms to express seemingly common designs such as node-link diagrams, parallel coordinates, tree diagrams, nested visualizations, nor tree maps.

For instance, Figure 1 connects the points in  $V_T$  based on the edges table  $E(id, s, t)$ . In existing approaches, the user first computes  $E' = T \bowtie E \bowtie T$ , then maps  $E'$  as links in  $V_{E'}$ . However, the edges only *appear* to connect the points. If the user jitters the points or applies a force-directed layout algorithm, then the positions of the points and edges become inconsistent because the visualization is unable to maintain their logical relationship (Figure 1(b)). Furthermore, single-table formalisms and libraries cannot support the force-directed layout algorithm because it needs to reason about *two* tables. As a consequence, the user needs to do more “preparation”: run the layout to compute pixel coordinates of the nodes in a new table  $T'$ , derive a new edges table  $E'' = T' \bowtie E \bowtie T'$ , and draw the points and links. This breaks the data-to-pixel mapping abstraction because  $T'$  and  $E''$  are tables of pixel coordinates, and the grammar is merely used for rendering because all of the layout logic was applied during preparation.

## 2.2 Relational Data Modeling

We denote a table by its schema  $T(a_1, \dots, a_n)$ , where attribute  $T.a_i$  has domain  $T.D_i$ , and use capitalization to denote a set of attributes  $T.X$ . We underline a key  $\underline{T.X}$ , and assume every table has primary key  $\underline{T.id}$ . A foreign key relationship  $C(S.X, T.Y)$  specifies that a subset of attributes  $S.X \subseteq S_S$  refer to attributes  $T.Y \subseteq S_T$  of another table<sup>1</sup>:  $\forall s \in S \exists t \in T, s.X = t.Y$ . 1-1 relationships are when both  $S.X$  and  $T.Y$  are keys, denoted  $C(\underline{S.X}, \underline{T.Y})$ ; N-1 relationships denoted  $C(S.X, \underline{T.Y})$ . N-M relationships are modeled using an intermediate relation  $W(X, Y)$  with constraints  $C(W.X, \underline{S.X})$  and  $C(W.Y, \underline{T.Y})$ .

## 2.3 Single-Table Grammars

We now introduce a simple formalism to describe graphical grammars that visually map an input table  $T$  to an output view  $V$  (a *visual mapping*). The notation focuses on data attributes, marks, mappings between data and mark properties, and scales. A visualization may have multiple views, often laid out next to each other, or super-imposed as layers.

In database parlance, the view  $V$  is a projection of data attributes to mark properties (often called a visual channel) that is indexed by the table’s keys. The notation  $\xrightarrow{s} v$  specifies an *aesthetic mapping* from data attribute  $a$  to mark property  $v$  using a scale function  $s$ . Named scale functions can be referenced in multiple views; by default each mapping uses

<sup>1</sup> Codd’s definition [3] is more strict, and states that  $S.X$  is not a key and  $T.Y$  is a primary key. We relax this definition in order to model relationships as used in practice.

a different scale function. For instance, the following creates a scatterplot that uses linear scales to map values of  $a$  to pixel positions from 10 to 100 along the  $x$  axis (similarly for  $b \xrightarrow{s_y} y$ , and the primary key  $id$  using an identity function ( $s_{id}$ )).

$$V = \{T \rightarrow \text{point}, T.\text{id} \xrightarrow{s_{id}} \text{id}, T.a \xrightarrow{s_x} x, T.b \xrightarrow{s_y} y\} \quad s_{id} = \text{identity}(), s_x = \text{linear}(), s_y = \text{linear}()$$

$V$  is a table of marks with schema  $V(\text{type}, \underline{\text{id}}, x, y, \dots)$  that contains its mark type, keys, and mark properties.  $V$  has a 1-1 relationship with  $T$  via the constraint  $C(\underline{V.\text{id}}, \underline{T.\text{id}})$ , so it preserves the same relationships that  $T$  has with other tables. We will use a simplified notation that omits the primary key mapping, table prefixes, and scale functions, unless they are important to an example. For instance, the following is equivalent to the above example:  $V = \{T \rightarrow \text{point}, a \rightarrow x, b \rightarrow y\}$ .

### 3 Database Visualization

*Database visualization* maps database contents *and constraints* to a visual representation. We describe single-table grammars as constraint-preserving mapping, introduce extensions to foreign-key constraints, and illustrate the connection between data modeling and visualization design.

#### 3.1 Graphical Grammars as Mapping Constraints

We first examine the existing single-table formalism from the perspective of preserving attribute domain and key constraints. Attribute domains are preserved by drawing axes or legends to illustrate how the data attribute's domain maps to the range of the visual channel. Without these visual metadata, marks are simply objects floating in space, and the user cannot interpret the visually encoded data.

To understand key constraints, let us revisit the scatterplot in the previous section:  $V = \{T \rightarrow \text{point}, \text{id} \rightarrow \text{id}, \text{year} \rightarrow x, \text{price} \rightarrow y\}$ . To preserve the key constraint  $\underline{T.\text{id}}$ , we must be able to distinguish different rows in the visualization. Is the constraint trivially satisfied because  $\underline{V.\text{id}} = \underline{T.\text{id}}$ ? Unfortunately no, because the visualization is not a set or bag of rows in an  $N$ -dimensional space – its marks are spatially positioned in a 2D space. Thus, different marks may appear indistinguishable and thus violate the constraint. This is commonly called “overplotting”.

What are ways to address this constraint violation? One common approach is to change the constraints by e.g., grouping by the spatial attributes ( $T.(\text{year}, \text{price})$ ) to render a heatmap where the spatial attributes are the key, or reducing opacity achieves a similar goal but groups at the pixel granularity. A second approach is to change the specification by e.g., mapping  $\text{id} \rightarrow x$  and  $\text{year} \rightarrow \text{color}$  to remove overlaps. The third is to perturb the layout by e.g., jittering the points to reduce overlaps. This variety of interventions is possible because the violations are visual and low dimensional. This contrasts with data cleaning, which resolves violations by changing data values.

#### 3.2 Data Modeling and Visualization Through Examples

We define a *faithful database visualization* as a mapping where 1) each table maps to one or more views; in a given view, 2) each row maps to one mark and 3) each attribute maps to a mark property; and 4) each constraint is preserved in the views. Single-table mappings are sufficient to express (1-3), so we now shift focus to (4) using the series of examples in

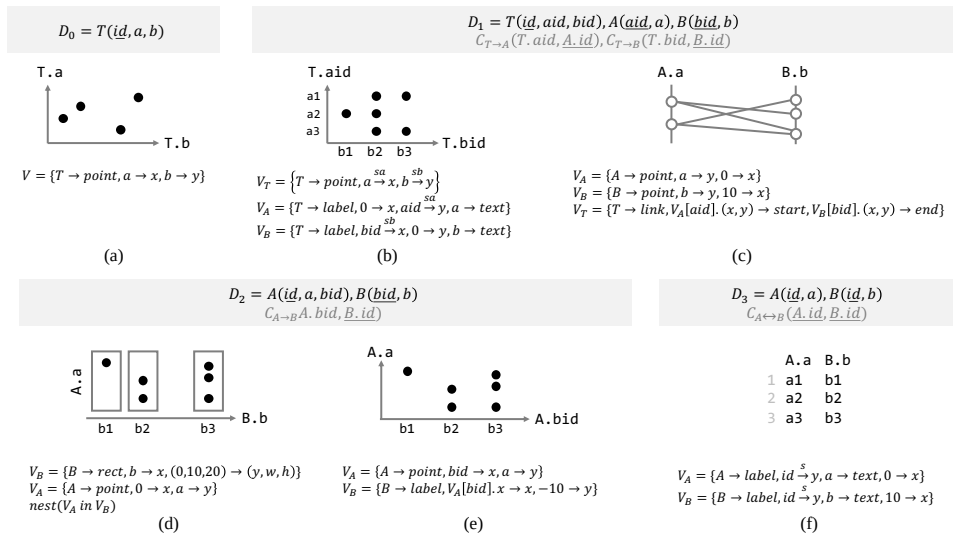


Figure 2 Examples that show how different visualization structures are needed to faithfully express different sets of tables and constraints. Common visualization designs are a consequence of data modeling decisions.

Figure 2. Starting with the database  $D_0$  in Figure 2(a), we different decompositions to illustrate how the data model affects the structure of the visualization, and several ways to visually preserve foreign keys.

### 3.2.1 Many-Many Decomposition

$D_1$  normalizes the attributes  $a$  and  $b$  as their own entities, and the original  $T$  encodes a many-to-many relationship between tables  $A$  and  $B$ , as is common for graph data. Figure 2(b) renders the cross product  $T.aid \times T.bid$  as the scatterplot  $V_T$  and marks each row in  $T$  with a point.  $A$  and  $B$  are rendered as labels along the  $y$  (in  $V_A$ ) and  $x$ -axes (in  $V_B$ ); the relationships are preserved by sharing the scales so that the labels are aligned with their positions in  $V_T$ . For instance  $V_A$  maps  $aid \rightarrow y$  using the same scale  $s_a$  as in  $V_T$  to preserve their correspondence.

Figure 2(c) visualizes the values of  $A.a$  and  $B.b$  along separate axes, and renders  $T$  as links.  $V_T$  does not directly map  $T.aid$  as in Figure 2(b). It instead uses  $T.aid$  to look up its corresponding point mark via  $V_A[aid]$ ; this follows the foreign key constraint  $C_{T \rightarrow A}$  and the 1 – 1 relationship from  $A$  to  $V_A$ , and then maps the point mark’s position (which is a key) to the link’s start position. Referencing  $V_A[aid]$  via the foreign keys allows them to be maintained in the visualization even as marks move. We note that this design, commonly called parallel coordinates, is a consequence of data modeling.

In general, the specification of a view of  $T$  may refer to any table  $T'$  for which there is an unambiguous join path  $T \rightsquigarrow T'$ , so that each tuple in  $T$  derives a unique value. In both of these examples, the foreign key relationships are preserved through alignment of scales or mark position.

### 3.2.2 Many-One Decomposition

$D_2$  only normalizes  $b$ , and is commonly used to express hierarchical data. Figure 2(d) visualizes  $A$  as a scatter plot  $V_A$  where the  $y$  position is determined by  $A.a$ , while  $B$  is visualized in  $V_B$  as rectangle marks whose  $x$ -positions are based on  $B.b$ . The visualization so

far is invalid because  $A.bid$  is not mapped and the constraint  $C(A.bid, B.bid)$  is not preserved. Since each  $B$  row maps to multiple  $A$  rows,  $nest(V_A \text{ in } V_B)$  preserves the constraint by drawing  $V_A$  within marks of  $V_B$ . Specifically, it treats each  $V_B$  mark as a “subcanvas” and renders  $V_A$  inside of it using its related rows in  $A$ . Depending on the individual view definitions, nesting expresses faceting, framed-rectangle plots, treemaps, and other nested visualizations.

Figure 2(e) renders  $A$  as a scatterplot where  $A.bid$  is mapped to the  $x$  position, and  $V_B$  renders the corresponding labels for  $B$ . This effectively treats  $b$  values as a categorical variable,

### 3.2.3 One-One Decomposition

The final decomposition  $D_3$  treats  $a$  and  $b$  as separate entities with a 1-1 relationship. Figure 2(f) maps  $A.a$  (and  $B.b$ ) to labels where their  $id$  determines their positions (similar to  $B.b$  in the previous visualization). Since  $V_A$  and  $V_B$  are aligned along the  $y$  position (derived from their input tables’ key), this visualization is faithful. This is commonly called a table.

## 4 Conclusion and Future Work

This short paper has outlined the basic principles for database visualization as a mapping from database constraints and contents to a faithful visual representation. Ultimately, database visualization defines a new design space that encapsulates many visualization designs, such as parallel coordinates and facets, that are not expressible under a single-table formalism. This work also raises questions of potential interest for database theorists and practitioners. To what extent do the lessons in data modeling – identifying salient entities and relationships, normal forms, cost estimation – translate to and explain visualization designs? How can layout algorithms (e.g. force-directed layouts) be incorporated into this formalism, and what is the line between declarative mapping specifications and imperative layout algorithms – to what extent are visualizations “just queries”, really? Is there a formalism that expresses the continuum between diagrams of the database metadata (e.g., ER diagram) and visualizations of the database contents – does this require a formalism over second-order queries? A common source of multi-table databases is the intermediate results generated by relational or data science pipelines. What constraints are present in the resulting database of intermediate results beyond those inferred by the chase [1]? What are the visual representations for constraints in non-relational data models such as event streams, matrices, and 3D representations?

---

### References

- 1 Alfred V Aho, Catriel Beeri, and Jeffrey D Ullman. The theory of joins in relational databases. *ACM Transactions on Database Systems (TODS)*, 4(3):297–314, 1979. doi:10.1145/320083.320091.
- 2 John J. Bertin. *The semiology of graphics*. University of Wisconsin Press, 1983. URL: <https://api.semanticscholar.org/CorpusID:59708204>.
- 3 Edgar F Codd. A relational model of data for large shared data banks. *Communications of the ACM*, 13(6):377–387, 1970. doi:10.1145/362384.362685.
- 4 Arvind Satyanarayan, Dominik Moritz, Kanit Wongsuphasawat, and Jeffrey Heer. Vega-lite: A grammar of interactive graphics. *IEEE Transactions on Visualization and Computer Graphics*, 23:341–350, 2018. doi:10.1109/TVCG.2016.2599030.

## 35:6 Database Theory in Action: Database Visualization

- 5 Chris Stolte, Diane Tang, and Pat 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. doi:10.1109/2945.981851.
- 6 Hadley Wickham. A layered grammar of graphics. *Journal of Computational and Graphical Statistics*, 19:28–3, 2010. URL: <https://api.semanticscholar.org/CorpusID:266499321>.
- 7 Leland Wilkinson. nvizn : An algebra-based visualization system. In *Smart Graphics*, 2001. URL: <https://api.semanticscholar.org/CorpusID:14945586>.
- 8 Leland Wilkinson. *The grammar of graphics*. Springer, 2012.