Presenting Visualization Guideline Collections

James Scott-Brown*
University of Oxford

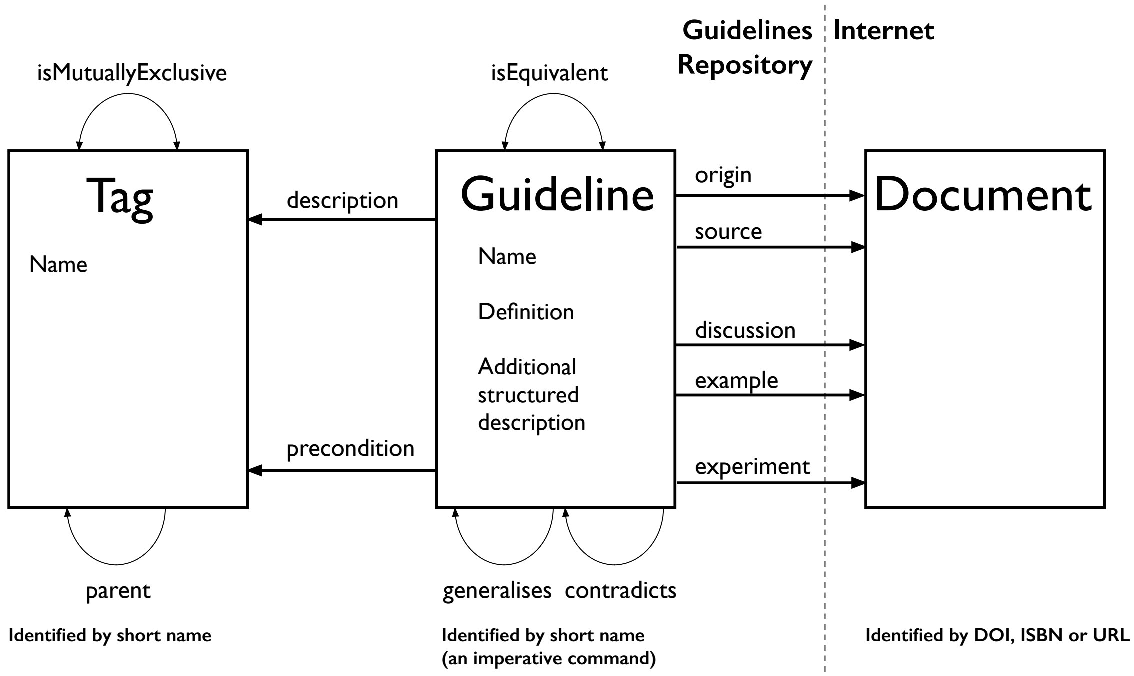

Figure 1: Proposed data model for a repository of guidelines.

Abstract

There have been a number of attempts to create collections of guidelines for visualizations, in several different formats. This position paper briefly reviews the structure of some of these attempts, suggests a general data model for recording a collection of guidelines, and describes some presentations that this would enable. It argues for a much more structured representation of guidelines than has been previously used.

1 Introduction

The most general repository of visualization guidelines would serve several different user needs, including acting as a *reference tool* to help users solve particular problems that they are currently encountering, as a *teaching tool* to help students learn general principles that they can apply to many different problems in the future, and as a *bibliography* and *record* of discussions that have occurred in the visualization community. Ideally it would be usable as a formalised knowledge base, parts of which could be fed into automated visualization design tools, but this is not all that it should be: it should also include additional contextual information, guidelines that are not expressed precisely enough to be used in automated reasoning, and a representation of contradictory advice given by different people.

*e-mail: james@jamesscottbrown.com
Present affiliation: Data Science Institute, Imperial College London.

2 Existing presentations of guidelines

The focus of this paper is on guidelines in visualization, but guidelines are also used in other scientific and technical fields. They are perhaps most prominently used in medicine, where they inform clinical practice based on the latest available evidence; given the volume of clinical research performed it would be difficult for clinicians to remain up to date from the primary literature alone. The Guidelines International Network coordinates between approximately 100 member organisations, and maintains an International Guideline Library containing over 6,000 guidelines [20]. A medical guideline is not a single atomic piece of advice, but rather a short monograph on a particular topic: for example, the single NICE *guideline* NG98 [28] is 282 pages long, but it contains multiple *recommendations* that are each a sentence long (e.g., "Offer to remove earwax for adults in primary care or community ear care services if the earwax is contributing to hearing loss or other symptoms, or needs to be removed in order to examine the ear or take an impression of the ear canal."): this is in contrast to usage in other fields (such as visualization), where each individual recommendation would be referred to as a guideline. Within the field of scientific research, the journal *PLOS Computational Biology* has published more than 90 editorials in its *Ten simple rules* format, covering a wide range of topics [30].

Collections of visualization guidelines are currently presented in a range of different formats, each of which is suited to different needs: including guidelines in continuous prose is helpful for teaching beginners, decision trees are helpful for guiding users through the initial choices about a visualization, and checklists are helpful when improving a draft visualization.

Text

Guidelines are often presented within expository text, such as a blog-post describing a single guideline, a paper that introduces a small number of visualizations, or a textbook that lists a large number of guidelines. The guidelines themselves have no particular structure, but rather fit into the structure of the surrounding text.

Decision Trees

The purpose of a guideline is to help the user to decide what they should do, and the most direct way to achieve this goal is with a decision tree that asks users a series of questions, then presents them with a suggested solution (or set of solutions) to their problem.

Several decision trees make suggestions of which types of chart to draw based upon task (e.g., comparison, distribution, composition, trend, relationship) and data-type. Both 'Chart Suggestions — A Thought Starter' [8] and 'from Data to Viz' [21] present a decision tree with a maximum depth of 4 choices as a diagram. Kopol.js [4, 31] is an interactive tool, rather than a static diagram: the user enters the task, the type of two data attributes, and an importance weighting for performance time, accuracy and popularity into a form, and then Kopol returns a ranking for 5 alternative chart types. A further step towards automation is taken by knowledge-based automated visualization design tools such as APT [25], the Compass recommendation engine inside Voyager [37] and Draco [26]. However, automation conceals the reason why particular recommendations were made, these systems do not explicitly present guidelines to the user.

Other chart chooser diagrams divide chart types based on only a single top-level choice, simply presenting an unordered set of suggestions for each choice [3, 16]. A similar role is served by visual survey browsers, such as those implemented using SurVis [13], which let a user filter a large set of visualization types based on multiple criteria. However, such survey browsers are typically targeted at visualization experts and include specialized visualizations, whereas the chart choosers listed above are aimed at helping inexperienced users choose between common chart types.

Decision trees have also been applied to colour-scheme selection: ColorBrewer2 [1] asks the number of data classes, whether the variable is sequential/diverging/qualitative, and whether the colourscheme needs to be colorblind safe/photocopy safe/print friendly, then presents a set of colour schemes consistent with this requirement.

The problem of choosing of map projection for geospatial data is substantive enough to be the subject of at least one dedicated monograph [24], in addition to the reference books that catalogue projections. Some resources [10] try to guide the user through the decision process by asking about the purpose of the map, the region of the world to be shown, and the shape of the area of interest, but this advice is not neatly structured into a single decision tree.

Lists

Checklists (e.g., [17]) provide a list of common faults, or areas for potential improvement, in visualizations. They are intended to be used as a guide when checking draft images.

A *visualization style guide* (e.g., [29, 33, 35]) may be written in a format similar to a checklist. The aim of such style guides is often to achieve consistency and 'create a uniform look and feel' [35] between different visualizations (and other documents) produced by the same organization; they mandate specific choices for details such as color, typeface, and line-width, and act partly as a quick reference for these arbitrary values. However, they do also include more general guidelines (e.g., 'Be sure to only use a map if the primary component of your data is geographical. Sometimes a ranked list or bar chart is better . . . ' [33]). Similar sets of guidelines are also used in software development, where *Human Interface Guidelines* provide software developers with recommendations about achieving a look and feel consistent with that of other applications on the same particular platform, so that new applications seem familiar and intuitive to a user familiar with other applications.

Pattern libraries

A *design pattern* is a reusable solution to a recurring design problem; a complete *pattern language* consists not only of a library of design patterns, but also a syntax describing how they can be combined. The term was introduced by the architect Christopher Alexander in his book *A Pattern Language* [11], which presented 253 architectural patterns, but has been widely applied in software (popularised by the 'Gang of Four' book [19], which presented 23 patterns grouped into 3 categories, and the original wiki [2]). The term 'idiom' is also used in some visualization texts with essentially the same meaning as pattern (e.g., [27]).

There have been several attempt to identify patterns for visualizations, but they have been limited in size and none have extended to form a pattern language. The `VizPatterns.org` wiki was merged into the `InfoVis:Wiki` in 2007 [9], but remains very incomplete. Colin Ware has collected 22 Visual Thinking Design Patterns for Data Analysis Tools as an unstructured list [36]. The NAPA Cards project includes 16 Narrative Patterns for Data-Driven Storytelling, each of which is assigned one or more of 5 possible tags ('empathy', 'engagement', 'framing', 'flow', and 'argument') [12].

Christian Behrens' created an interactive browser of Information Design Patterns that presented 48 patterns, divided into 8 categories of 'display patterns' (correlations, continuous quantities, discrete quantities, proportions, flows, hierarchies, networks, spatial configurations), 5 categories of 'behaviour patterns' (navigation, filtering, arrangement, exploration, transition), and 3 categories of 'interaction patterns' (discrete selection, linear adjustment, spatial navigation). The browser is no longer online, but the corresponding master's thesis is still accessible [14].

There have also been a number of attempts to create pattern libraries for Human Computer Interaction more generally [5]; [18] lists 30 such libraries, grouped according to the pattern format that they use.

However, a design pattern is not quite the same as a guideline. Pattern libraries are generally constructed by one person (or a small group of people) extracting patterns directly from what they consider to be successful designs; we propose instead aggregating the large number of existing guidelines in visualization into one structured resource containing guidelines that were initially proposed by many different people. This will result in a more complex structure, as authors may disagree and propose conflicting guidelines, whereas a pattern language should be consistent.

Statistics

The authors of [22] extracted 'about 550 guidelines' from 5 books, 18 papers, and 1 blog, but they presented only statistical information about the form and content of these guidelines, rather than the actual guidelines themselves. They performed a grounded theory analysis and constructed a three-level hierarchy of concepts, sub-concepts and instances contained in the guidelines.

3 Proposed Data Model

The proposed *Guideline Report* format for this workshop [6] lists details that should be included in a complete description of a single guideline. However, in constructing a collection of guidelines it is desirable to explicitly represent relationships between guidelines, and between guidelines and the external literature. A proposed data model to capture these relationships is shown in Figure 1.

One guideline may *generalise* another (e.g., 'remove unnecessary grid-lines, or reduce their visual weight' can be considered a special case of 'maximise the data ink ratio'), or *contradict* another (e.g., a recommendation to 'add decoration to increase memorability' might contradict a recommendation to 'avoid chart-junk').

A guideline may have a clear *origin* where it was first proposed: 'maximise the data-ink ratio' can be traced to [34], and Shneiderman's Visual Information Seeking Mantra to [32]. There is no single origin for guidelines such as 'avoid using pie-charts', but it would still be useful to record the reference from which the guideline was actually copied as a *source*.

A guideline may also be *mentioned* or *discussed* in many places: it is infeasible to collate all mention of common guidelines, but it may be useful to include references to particularly insightful discussions. However, it is feasible to automatically list all discussions on the VisGuides forum [7] that discuss a particular guideline, if they either include the name of a guideline or a link to its entry.

There are likely to be *example images* of visualizations that either do or do not comply with particular specifications; some of these will be included in discussions of guidelines, but others will occur 'in the wild'. Linking to these, rather than including them directly in the repository, may avoid the hassle of securing copyright permissions. The creator of a guideline entry could generate images for the specific purpose of illustrating the guideline, but this would involve considerable effort.

A guideline may be supported or rebutted by *experimental evidence*. Evidence may support a guideline without having directly tested it: for example, a study of the incidence of colour-blindness could be used as justification for the guideline 'use a colorblind-safe palette'. A subtlety is that a guideline is a suggestion of *what* to do, rather than being an explicit claim about performance: an empirical evaluation of a guideline most focus on a *specific effect* of following a guideline — for example, whether the guideline improves aesthetics, memorability, recall, reading time, or reading accuracy.

These links should refer to external entities using standard identifiers (Digital Object Identifiers for papers, International Standard Book Numbers for books, Uniform Resource Identifiers for blog posts and other online content). To facilitate linking into a repository of visualization guides, each guideline should be assigned a unique, stable identifier that does not depend upon the classification of the guideline into a hierarchy. In the interest of memorability, these could be given short names that are imperative commands (e.g., 'AvoidPieCharts', 'MaximiseDataInkRatio'); it is understood that whilst these names are commands that suggest a high degree of confidence, they should be viewed as suggestions and be critically interpreted rather than applied blindly.

3.1 Tags

Rather than trying to structure guidelines into a hierarchy directly, it may be better to arrange *tags* into a hierarchy, and associate each guideline with zero or more tags, allowing guidelines to be categorized along multiple independent axes.

A qualitative distinction can be made between tags that *describe a guideline itself*, and tags that *describe the conditions required for a guideline to apply*. For example, description tags could indicate:

- the *general aspect of a visualization* to which the guideline applies, such as the use of **color**, **interactivity**, or **animation**.
- an assessment of how *empirically justified* a rule is
- a specific *decision* that a guideline informs (e.g., **Chart Type**, **Color Scheme**, or **Map Projection**). The child tags of such a tag would represent the alternative choices that must be decided between.

Precondition tags could indicate:

- the *chart type* to which the guideline applies: whilst some guidelines apply to all chart-types, others are specific to a single chart type or types (e.g., 'ensure totals sum to 100%' for **pie charts** or **stacked bar charts**)
- the *number of variables/dimensions* represented
- the *data type* for each variable
- the *task type*, or what the visualization is primarily trying to show (e.g., **comparison**, **distribution**, **relationship**, **composition/part-to-whole**)
- whether the purpose of the visualization is **exploration** or **communication/presentation**

Each of the items included in the two lists above could be a top-level entry in a hierarchy of tags. These tag hierarchies do not need to be constructed from scratch, as there are existing taxonomies that could be incorporated: for example, there are many existing task classifications [15, 23]

A pair of precondition tags may be mutually exclusive (for example **Variable Count** → **Two** and **Variable Count** → **Three**); this can be recorded as a *mutuallyExclusive* relation. Whilst such relations would typically be obvious to a human reader, representing them explicitly in the repository allows reasoning about them as described below.

A guideline might apply when some Boolean expression of conditions is satisfied (e.g., 'A AND (B OR C)'). This could be captured by interpreting a guideline as applying when all of the associated preconditions are satisfied, re-writing such expressions in Disjunctive Normal Form (e.g., '(A AND B) OR (A AND C)') and duplicating guidelines as necessary, and introducing an *isEquivalent* relation to record that the duplicated guidelines are equivalent.

As a concrete example of how a guideline expressing a decision could be encoded, suppose we want to encode the guidance expressed in [8] about when to use a scatter chart. We can interpret the decision tree as stating: 'A scatter plot is an appropriate representation for the relationship between two variables, or the distribution of two variables'. We can split this into two guidelines, 'A scatter plot is an appropriate representation for the relationship between two variables' and 'A scatter plot is an appropriate representation for the distribution of two variables', recording [8] as the *source* for both. Both could be labeled with the *description tags* **Decision** → **Chart-Type** (representing the question that they help decide) and **Decision** → **Chart-Type** → **Scatter Plot** (representing the choice that they suggest), and the *precondition tags* **Variable Count** → **Two, Purpose** → **Relationship** and **Variable Count** → **Two, Purpose** → **Distribution** (representing the conditions under which they make this suggestion). Some guidelines may suggest several alternative options; this could be encoded by associating the guideline with the tag representing each suggested choice, and recording an ordinal ranking for each.

4 Proposed guideline presentations

The proposed data-model would provide a rich representation of guidelines and their relationships, allowing a range of presentations.

An individual page for each guideline with a canonical, stable URI would act as a single point of reference, providing a complete description and links to related guidelines, tags, external documents (e.g., sources, examples, discussions, experimental evidence), and discussions on the VisGuides forum [7].

At the opposite extreme of granularity, the complete set of guidelines could be presented as a list or table that is filterable based on tags and substrings of guideline names or descriptions. A filtered list of guidelines could be exported as a checklist, including the name and a brief description of each.

The tag system could be presented as a hierarchy, with the ability to navigate to a page for each tag that lists the corresponding guidelines. The number of guidelines associated with each tag could be represented as a treemap.

From guidelines encoding decision rules, it is possible to automatically generate a decision tree for a particular decision. For example, from an encoding of the guidelines expressed in [8]), combined with the *mutuallyExclusive* relations for preconditions, it is possible to construct a decision tree for choosing between the alternatives represented by the child tags of the **Decision** → **Chart-Type** tag. Such decision trees could be displayed in a number of visual formats, such as an indented list, a flow-chart, or a matrix. Alternatively, a user could simply filter the list of guidelines by first selecting only guidelines with the **Decision** → **Chart-Type** tag, and then successively narrowing this selection by adding more precondition tags describing their data until they arrive at a recommendation, without explicitly reconstructing a decision tree[1]. If the meaning of tags representing preconditions and choices were formally defined, it would in principle be possible to export a selection of guidelines as a set of constraints that could be used with an automated visualization tool such as Draco [26].

The relationships between guidelines and documents could be visualized as a bipartite graph.

5 Conclusion

Previous collections of guidelines have been limited in scope. We envision the creation of a manually curated collection that is more ambitious in the number of guidelines included, the amount of detail represented about each guideline, and the flexibility of exploration.

References

[1] Colorbrewer 2.0 — color advice for cartography. `http://colorbrewer2.org/`. Accessed: 2018-08-23.

[2] Content creation wiki. `http://wiki.c2.com/`. Accessed: 2018-08-23.

[3] The Financial Times Visual Vocabulary. `http://ft.com/vocabulary`. Accessed: 2018-08-23.

[4] kopol.js visualization recommender. `https://kopoljs.github.io/`. Accessed: 2018-08-23.

[5] `hcipatterns.org` : patterns. `http://www.hcipatterns.org/patterns`. Accessed: 2018-08-23.

[6] VisGuides: 2nd workshop on the creation, curation, critique and conditioning of principles and guidelines in visualization. `http://aviz.fr/~bbach/visguides/`. Accessed: 2018-08-23.

[7] VisGuides Forum. `http://visguides.org`. Accessed: 2018-08-23.

[8] A. Abela. Chart suggestions - a thought starter. `http://extremepresentation.typepad.com/files/choosing-a-good-chart-09.pdf`. Accessed: 2018-08-23.

[9] W. Aigner. 2007-10-18: InfoVis:Wiki and VizPatterns Merger. `https://infovis-wiki.net/wiki/2007-10-18:_InfoVis:Wiki_and_VizPatterns_Merger`. Accessed: 2018-08-23.

[10] J. Albrecht. Choosing a projection. `http://www.geo.hunter.cuny.edu/~jochen/gtech201/lectures/lec6concepts/map%20coordinate%20systems/how%20to%20choose%20a%20projection.htm`. Accessed: 2018-08-23.

[11] C. Alexander. *A pattern language: towns, buildings, construction*. Oxford University Press, New York, 1977.

[12] L. Bartram, J. Boy, P. Ciuccarelli, S. Drucker, Y. Engelhardt, U. Koeppen, M. Stefaner, B. Tversky, and J. Wood. NAPA cards preview: Narrative patterns for data-driven storytelling. `http://napa-cards.net/`. Accessed: 2018-08-23.

[13] F. Beck, S. Koch, and D. Weiskopf. Visual analysis and dissemination of scientific literature collections with SurVis. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):180–189, jan 2016. doi: 10.1109/tvcg.2015.2467757

[14] C. Behrens. The form of facts and figures. `https://christianbehrens.com/portfolio/project/foff`, 2008. Accessed: 2018-08-23.

[15] M. Brehmer and T. Munzner. A multi-level typology of abstract visualization tasks. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2376–2385, dec 2013. doi: 10.1109/tvcg.2013.124

[16] A. K. Emery. Emery's essentials. `http://annkemery.com/essentials/`. Accessed: 2018-08-23.

[17] S. Evergreen and A. K. Emery. Data visualization checklist. `http://stephanieevergreen.com/interactive-data-visualization-checklist/`. Accessed: 2018-08-23.

[18] S. Fincher. The pattern gallery. `https://www.cs.kent.ac.uk/people/staff/saf/patterns/gallery.html`. Accessed: 2018-08-23.

[19] E. Gamma. *Design patterns : elements of reusable object-oriented software*. Addison-Wesley, Reading, Mass, 1995.

[20] Guidelines International Network. International guideline library. `https://www.g-i-n.net/library/international-guidelines-library`. Accessed: 2018-08-23.

[21] Y. Holtz and C. Healy. Data to viz. `https://www.data-to-viz.com/`. Accessed: 2018-08-23.

[22] E. Kandogan and H. Lee. A grounded theory study on the language of data visualization principles and guidelines. *Electronic Imaging*, 2016(16):1–9, feb 2016. doi: 10.2352/issn.2470-1173.2016.16.hvei-132

[23] N. Kerracher and J. Kennedy. Constructing and evaluating visualisation task classifications: Process and considerations. *Computer Graphics Forum*, 36(3):47–59, jun 2017. doi: 10.1111/cgf.13167

[24] M. Lapaine and E. L. Usery, eds. *Choosing a Map Projection*. Springer International Publishing, 2017. doi: 10.1007/978-3-319-51835-0

[25] J. Mackinlay. Automating the design of graphical presentations of relational information. *ACM Transactions on Graphics*, 5(2):110–141, apr 1986. doi: 10.1145/22949.22950

[26] 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 Trans. Visualization & Comp. Graphics (Proc. InfoVis)*, 2018.

[27] T. Munzner. *Visualization analysis & design*. CRC Press/Taylor & Francis Group, Boca Raton, FL, 2015.

[28] National Institute for Health and Care Excellence. Hearing loss in adults: assessment and management (NG98), 2018. Accessed: 2018-08-23.

[29] Office for National Statistics. Data visualisation - Style.ONS. `https://style.ons.gov.uk/category/data-visualisation/`. Accessed: 2018-08-23.

[30] PLOS Computational Biology. PLOS collections: 10 simple rules. `http://collections.plos.org/ten-simple-rules`. Accessed: 2018-08-23.

[31] B. Saket, A. Endert, and C. Demiralp. Task-based effectiveness of basic visualizations. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–1, 2018. doi: 10.1109/tvcg.2018.2829750

[32] B. Shneiderman. The eyes have it: a task by data type taxonomy for information visualizations. In *Proceedings 1996 IEEE Symposium on Visual Languages*. doi: 10.1109/vl.1996.545307

[33] Sunlight Foundation. Data visualization style guidelines. `https://github.com/amycesal/dataviz-style-guide/blob/master/Sunlight-StyleGuide-DataViz.pdf`. Accessed: 2018-08-23.

[34] E. Tufte. *The visual display of quantitative information*. Graphics Press, 1983.

[35] Urban Institute. Urban institute data visualization style guide. `http://urbaninstitute.github.io/graphics-styleguide`. Accessed: 2018-08-23.

[36] C. Ware. Visual thinking design patterns for data analysis tools. `https://ccom.unh.edu/vislab/VTDP_web_pages/VTDP_HomePage.html`. Accessed: 2018-08-23.

[37] K. Wongsuphasawat, D. Moritz, A. Anand, J. Mackinlay, B. Howe, and J. Heer. Voyager: Exploratory analysis via faceted browsing of visualization recommendations. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):649–658, jan 2016. doi: 10.1109/tvcg.2015.2467191

[1]This would require a filtering option that selects not only guidelines that have a particular tag, but also guidelines that do not have it but have no other tags that are *mutuallyExclusive* with it.