Tutorial Proposal: Connecting the Dots – Showing Relationships in Data and Beyond

Marc Streit*
Johannes Kepler University Linz

Hans-Jörg Schulz†
University of Rostock

Alexander Lex‡
Harvard School of Engineering and Applied Sciences
Graz University of Technology

ABSTRACT

Relationships are omnipresent in data, views and in how we interact with visualization tools. This tutorial discusses how such relationships can be visually expressed, a process we call **linking**. The tutorial addresses the three questions, **what**, **how** and **when** to link in three separate parts. The first part – what to link – explains that not only data, but also views and interactions can be linked. The second part – how to link – discusses how relationships can be visually expressed, based on perceptual grouping principles. While we discuss a wide range of methods, we focus on the *connectedness* grouping principle, specifically on *visual links*, as the topic is the most powerful in some respects, but also the most difficult to employ. The final part – when to link – will give an introduction to scenarios where linking is beneficial, taking into account issues such as unconventional display devices and collaboration. Again the focus will be on visual links, but other techniques will be discussed where appropriate.

1 TUTORIAL CONTENTS

Comparing, evaluating and interpreting related pieces of information are fundamental tasks in visual data analysis but also in any kind of information intensive work in general. This tutorial for discusses the utility and design considerations of visually expressing relationships between such pieces of information. We refer to this process of expressing relationships as “linking”. The three parst of the tutorial are the following:

1. **Part 1: What to link?**
   Relationships can be contained in the data, can exist between views, as well as on the interaction level.
2. **Part 2: How to link?**
   This part is concerned with the different possibilities for visually expressing relationships.
3. **Part 3: When to link?**
   In the third parts we discuss scenarios for visual linking, such as linking for multiple users, across applications, in unconventional display environments, etc.

1.1 Part 1: What to link?

The first part of the tutorial will elaborate on the different levels on which relationships can exist in the context of visualization. These are:

Data Level Relationships Relationships in the data can be defined by a similarity or dependency in the data. Relationships can be 1:1 but also 1:n and n:m. These relationships can exist on any level of data granularity:

1. between attribute values of a data item, such as belonging to the same data tuple – a relation that is regularly used, *e.g.,* in parallel coordinates,
2. between multiple data items, as determined, for example, by clustering algorithms,
3. between groups of data items, *e.g.,* when sharing the same data items [24, 22],
4. between entire datasets, *e.g.,* when there are foreign key relations between them [42],
5. between groups of items of different datasets, basically as a combination of 3 and 4 [25]

View Level Relationships Relationships on the view level capture dependencies between different parts of a visual representation. This is mostly done when views are displaced and shown at a different location. For example, it is quite common that an overview indicates the relation to an adjoint detail view by showing which part of it is displayed in detail [16]. Another instance, in which displacement is communicated to a user is when geospatial information is repositioned, such as, for example by Wood and Dykes [50].

Interaction Level Relationships Data items and views can also stand in a relationship due to their combined use. Such can be determined either from a preconceived workflow [42] or, for example, from logged interaction as in VisTrails [1]. Such relationships are usually communicated in a separate view that shows the dependencies between data and views to be used in concert in a certain order, such as Kreuseler's History Tree [20] or Aruvi's navigation view [39]. In some cases, such information can also be found integrated into the UI elements, such as done by Scented Widgets [48].

This part of the tutorial will elaborate on each of these sources of relationships. It will detail for each individually, with which methods to capture and evaluate such relationships, including clustering fundamentals and different concepts to log visualization sessions.

1.2 Part 2: How to link?

After discussing the different levels on which relationships in datasets can occur, part two will be dedicated to the different means of representing them. We distinguish different classes of links based on Gestalt principles [47] and recent extensions [31, 30] that constitute perceptual grouping principles. For an excellent review of perceptual grouping and other perceptual consideration in visualization refer to Healey's and Enns' review [9] or to the accompanying website[1].

- **Similarity** – A visual variable other then position is co-modulated between linked entities or (identical or related) glyphs (*e.g.,* a pointer or a label) are added in immediate proximity or connected to the linked items. Alternatively, everything but the linked entities is modulated, *e.g.,* by darkening or blurring.

*e-mail: marc.streit@jku.at
†e-mail: hjschulz@informatik.uni-rostock.de
‡e-mail: alex@seas.harvard.edu

[1] http://www.csc.ncsu.edu/faculty/healey/PP/index.html

- **Proximity** – Linked entities are placed in close proximity.
- **Connectedness** and **common region** – Linked entities are surrounded or connected by a geometry.

Other perceptual grouping principles, such as **common fate**, **good continuity** and **past experience** and **good continuation** will be briefly mentioned, but their applications are rare for various practical reasons.

While we discuss representatives of all classes, we focus on those that employ connectedness, since connectedness is both a powerful grouping principle but also sometimes difficult to employ, as, for example, its extensive use can lead to clutter. Also, there has been much research in the area lately with many novel techniques presented.

Similarity Color similarity is being almost universally employed for highlighting and brushing in visualization systems. Most Multiple Coordinated View (MCV) systems, like, for example, *Tableau* [41], or *VisPlore* [33], employ colored brushing to connect views such as parallel coordinates, histograms, or scatterplots. Color was also used in the when Linking+Brushing was first introduced [2]. Reasons for the widespread adoption of color are its ease of implementation, its ability to concurrently highlight arbitrary numbers, and its preattentive properties [44, 49]. Preattentive entities are recognized immediately, independent of the number of distractors. Non-preattentive attributes require serial search, *i.e.,*, conscious (attentive) comparison of every item. Employing color for highlighting, however, also has several drawbacks. While color is ideally suited to encode many items of one class, color is ill suited to encode many classes, *i.e.,*, the selective properties of color are limited. Healy [10] found that more than seven colors lead to reduced performance in accurately and rapidly detecting the colors. Also, color may be already employed to encode other attributes. SimVis [5], for example, uses color to encode other parameters, and falls back to (binary) saturation to highlight brushed areas. Equally, color can not be used for highlighting when color is used to encode a value, such as in pixel-based techniques.

Synchronous blinking is another form of similarity. While it is preattentive, it is also considered disturbing by many users and can hardly be used for more than one or two items. Techniques that add a symbol or glyph, for example a frame or an underline, or a label are another method to make use of similarity. Drawing glyphs or symbols and labeling can theoretically encode many relationships simultaneously. However, glyphs and labeling are even less selective than color, meaning that finding two related items requires serial search when enough distractors are present.

Modulation of content can also be used inversely, meaning that only the portions of a view that are to be linked are not modulated, and therefore the "not-modulation" is what constitutes the similarity. Modulation of the surrounding is typically done by decreasing saturation [51], brightness [51, 17], or sharpness [18]. Zhai *et al.* [51] show that darkening and decreasing saturation are highly effective but negatively affect user satisfaction. Hoffmann *et al.* [11] reproduced the negative user rating for darkening, and found that darkening was more error-prone than colored highlighting or connectedness. Kosara *et al.* [19] found that blurring is also highly effective as a highlighting technique. We can generalize that the modulation of the surrounding is very effective, but not versatile and not scalable. In fact, it is not possible to express more than one relationship at once. Also, the implementation is sometimes not straight-forward (blurring may require shaders, for example). Combined with the low user satisfaction this may be the reason why these techniques are not widely used for highlighting (synchronous or individual), even though all these techniques are preattentive.

Proximity Proximity as a grouping/linking principle requires that the visual variable position is not used to encode data. Consequently, interactive linking using proximity is mostly used in graph visualization, where users can select and drag portions of graphs to a shared location to symbolize a relationship (*e.g.,* [26]). Sorting or filtering of tables based on specific criteria (*e.g.,* [34, 38]) arranges items with similar or shared features in close proximity to each other as well. Automatically derived relationships are often encoded using position, for example in multidimensional scaling approaches [43], or in projections of word clouds based on semantics [32].

Connectedness Connectedness (or *uniform connectedness* [30]) was shown to be a very strong grouping principle, even stronger than classic *Gestalt* principles [47] such as proximity, similarity, or common fate [30, 52]. It was also shown that connected elements are perceived preattentively, approximately at the same speed as proximity, but faster than similarity [8]. Ziemkiewicz and Kosara distinguish between three forms of connectedness, namely connector, outline, and fill [52]. The latter two are what Palmer refers to as *common region* [31].

We distinguish between **general links**, as they are, for example, used in node-link diagrams (where the links are representation of the edges, which are part of the data structure), or in parallel coordinates plots (where the links encode the actual information) and visual links. We define **visual links** as "continuous shapes such as connection lines, curves, or surfaces that connect or surround multiple related pieces of information, thereby augmenting a base representation" [40]. In this context, a base representation is a image or visualization that is meaningful without the addition of visual links. The notion of base representation sets visual links apart from the general links, as they are used in node-link diagrams, where the meaning of the diagram is lost if the links are not present. There are two types of base representations. The first one is not aware of or does not adapt to visual links at all, *i.e.,*, visual links are superimposed on existing visualizations. The second type of base representations may leave empty space for the visual links, or may adjust its content for improved links routing.

The power of Visual Links is their unique ability to encode relations stronger than other methods, such as proximity, color, size or shape, can [30, 52]. Steinberger *et al.* [40], for example, have shown that search tasks can be performed more efficiently with visual links compared to colored highlighting. Consequently, Visual Links have gained widespread attention in the visualization community in recent years, with many publications on routing (*e.g.,* [7, 12, 13]), on their application to multiple simultaneous visualizations (*e.g.,* [3, 23, 45]), and on their use as a general tool to show and encode relations (*e.g.,* [4, 21, 24, 27, 35]).

Linking entities, albeit being a strong grouping and highlighting principle, does not scale well. As the number of links increases, their paths become hard to follow. Bundling strategies have been developed to group and organize the links and make linking scale to larger numbers. Bundling strategies either utilize an underlying structure, such as a hierarchy [12, 14]; use a force-directed layout where links attract each other [13]; or formulate the problem as an optimization to minimize the required ink [7]. A related problem, which is more relevant for visual links than general links, is that clutter makes the underlying base representation hard to read, an issue which is addressed using context-preserving visual links [40].

1.3 Part 3: When to link?

Having discussed where the relationships can come from and how they can be visually expressed, part three will elaborate on the different application scenarios that can potentially benefit from linking, specifically visual linking. We distinguish linking scenarios according to four variables – whether the setup comprises one or multiple **visualizations**, **applications**, **displays** and/or **users**. All kinds of combinations between these variables can be found in state-of-the-art systems and will be demonstrated by means of concrete examples from the literature. The higher the heterogeneity of the setup is with respect to these variables, the more challenging it is to come up with an effective linking solution.

Linking within a single visualization Linking is often used to encode additional relationships in a single visualization. Examples that employ visual links are simple node-link diagrams with hyper-edges added on top, but also Fekete's treemap overlays[6], Hierarchical Edge Bundles [12] and ArcTrees [28].

Linking across multiple visualizations Relationships between multiple visualizations are the most common linking scenario in visualization systems. Basically all MCV systems fall into this category. There are also various examples of applications using visual links between multiple visualizations. In the tutorial we will use the formalism proposed by Collins and Carpendale in their VisLinks work [3] to classify existing systems with respect to the interplay between datasets, relations within datasets and visualizations. However, their formalism stops at single applications. Therefore, in the final part of the tutorial we extend the formalism to also cover scenarios that include multiple applications, multiple displays as well as multiple users.

Linking across multiple applications Many state-of-the-art visualization systems allow users to analyze their data in a multiple coordinated view fashion. However, creating one single super-application that supports users in all their analysis needs is unrealistic. In real world situations, a profound visual analysis often comprises multiple, highly specialized and expensive tools. We are going to discuss approaches that allow to visually express relationships across multiple applications, for instance the Snap-Together system [29] and our own work on visual links across applications [45].

Linking across multiple displays It lies in the nature of visualization that the available number of pixels is a restricting factor. Even for medium-scale data analysis scenarios, the number of data items exceeds the available pixels for visualizing them. Abstraction techniques are one way to handle the problem, another way is to extend the available screen real estate. The spectrum ranges from a high resolution multi-projector system to off-the-shelf multi-desktop setups. In such setups the distance between the elements to link naturally increases, possibly even placing relevant entities out of a user's field of view. We will discuss the design choices made in a series of well established multi-display environments.

Linking for multiple users Collaboration is known to be an important building block for solving complex domain problems. For a comprehensive analysis, experts from multiple domains with different background knowledge are beneficial. However, in such multi-user scenarios the challenge lies in the coordination of the involved users. We will conclude part three by again presenting selected work on multi-user systems where linking plays a role, for instance, Isenberg and Fisher's collaborative linking and brushing [15] as well as Waldner's work on collaborative visual linking [46].

2 Instructor Information

Marc Streit is assistant professor at the Johannes Kepler University Linz in Austria and currently also a visiting researcher at the Center for Biomedical Informatics at Harvard Medical School. He received his PhD from Graz University of Technology in 2011. He has won the Best Paper Award at InfoVis'11 and GI'10, and the 3rd Best Paper Award at EuroVis'12. In 2013 he will be co-editing the IEEE Computer Special Issue on Visual Analytics. His research focuses on information visualization and visual analytics, where he is particularly interested in topics such as the analysis of heterogeneous datasets as well as visual linking. His research is embedded in the Caleydo project [2].

[2]http://www.caleydo.org

Alexander Lex is a post-doctoral researcher at the Visual Computing Group at Harvard School of Engineering and Applied Sciences. He received his PhD in 2012 from Graz University of Technology in Austria. His research interests are visualization especially in the context of molecular biology and human-computer interaction. He focuses on divide and conquer visualization for multidimensional inhomogeneous our heterogeneous datasets, on pathway visualization as well as on visual linking. He is one of the core members of the Caleydo project.

Together Marc and Alex have so far published numerous papers directly related to visual linking, among them one paper dealing with view arrangement for efficient visual linking in 2.5D and 2D scenes [23], one employing visual linking to show relations between multiple clustered sub-groups of multi-dimensional data sets [24], one discussing how to link across applications [45] and one paper on context preserving visual links [40]. The last two articles have won the *best paper award* at GI'10 and InfoVis'11 respectively.

Hans-Jörg Schulz received his PhD in 2010 from the University of Rostock, where he is now a post-doctoral researcher and also an associated researcher at the Graz University of Technology. His research interests include graph visualization and exploration, as well as their application to heterogeneous information landscapes, containing multiple interlinked data sets from various sources. He often employs his research in the context of the biomedical domain and systems biology. With his expertise in graph visualization, his perspective on showing relationships in data stems from a graph drawing point of view. While his main interest lies with visualizing hierarchical relationships in data [37, 36], he also, for example, utilized visual links and link-based interaction techniques for bipartite graph visualization [38]. As a self-confessed tree visualization aficionado, he spends his free time collecting entries for his Visual Bibliography of Tree Visualization (http://treevis.net).

The presenters have recently published an article describing how visual linking can be used in an heterogeneous, guided analysis scenario [42] as well as two papers where linking is heavily used to connect multi-form visualizations [22, 25]. The latter has won the *3rd best paper award* at EuroVis'12.

References

[1] L. Bavoil, S. Callahan, C. Scheidegger, H. Vo, P. Crossno, C. Silva, and J. Freire. VisTrails: enabling interactive multiple-view visualizations. In *Proceedings of the IEEE Conference on Visualization (VIS '05)*, pages 135–142. IEEE Computer Society Press, 2005.

[2] R. A. Becker and W. S. Cleveland. Brushing scatterplots. *Technometrics*, 29(2):127–142, May 1987.

[3] C. Collins and S. Carpendale. VisLink: revealing relationships amongst visualizations. *IEEE Transactions on Visualization and Computer Graphics (InfoVis '07)*, 13(6):1192–1199, 2007.

[4] C. Collins, G. Penn, and S. Carpendale. Bubble sets: Revealing set relations with isocontours over existing visualizations. *IEEE Transactions on Visualization and Computer Graphics (InfoVis '09)*, 15(6):1009–1016, 2009.

[5] H. Doleisch. SIMVIS: interactive visual analysis of large and time-dependent 3D simulation data. In *Proceedings of the Conference on Winter Simulation (WSC '07)*, page 712720, Piscataway, NJ, USA, 2007. IEEE Press.

[6] J.-D. Fekete, D. Wang, N. Dang, A. Aris, and C. Plaisant. Interactive poster: Overlaying graph links on treemaps. In *Proceedings of the IEEE Symposium on Information Visualization Conference Compendium (InfoVis '03)*, pages 82–83. IEEE Computer Society Press, 2003.

[7] E. R. Gansner, Y. Hu, S. North, and S. Carlos. Multilevel agglomerative edge bundling for visualizing large graphs. In *Proceedings of the IEEE Symposium on Pacific Visualization (PacificVis '11)*, pages 187–194. IEEE Computer Society Press, 2011.

[8] S. Han, G. W. Humphreys, and L. Chen. Uniform connectedness and classical gestalt principles of perceptual grouping. *Perception & Psychophysics*, 61(4):661–674, Jan. 1999.

[9] C. Healey and J. Enns. Attention and visual memory in visualization and computer graphics. *IEEE Transactions on Visualization and Computer Graphics*, 18(7):1170 –1188, July 2012.

[10] C. G. Healey. Choosing effective colours for data visualization. In *Proceedings of the IEEE Conference on Visualization (Vis '96)*, pages 263–270. IEEE Computer Society Press, 1996.

[11] R. Hoffmann, P. Baudisch, and D. S. Weld. Evaluating visual cues for window switching on large screens. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*, pages 929–938, 2008.

[12] D. Holten. Hierarchical edge bundles: Visualization of adjacency relations in hierarchical data. *IEEE Transactions on Visualization and Computer Graphics (InfoVis '06)*, 12(5):741–748, 2006.

[13] D. Holten and J. van Wijk. Force-directed edge bundling for graph visualization. *Computer Graphics Forum (EuroVis '09)*, 28(3):983–990, 2009.

[14] D. Holten and J. J. v. Wijk. Visual comparison of hierarchically organized data. *Computer Graphics Forum (EuroVis '08)*, 27:759–766, 2008.

[15] P. Isenberg and D. Fisher. Collaborative brushing and linking for colocated visual analytics of document collections. *Computer Graphics Forum (EuroVis '09)*, 28(3):1031–1038, 2009.

[16] W. Javed and N. Elmqvist. Stack zooming for multi-focus interaction in time-series data visualization. In *Proceedings of the IEEE Symposium on Pacific Visualization (PacificVis '10)*, pages 33–40, 2010.

[17] A. Khan, J. Matejka, G. Fitzmaurice, and G. Kurtenbach. Spotlight: directing users' attention on large displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '05)*, page 791798. ACM Press, 2005.

[18] R. Kosara, S. Miksch, and H. Hauser. Semantic depth of field. In *Proceedings of the IEEE Symposium on Information Visualization (InfoVis '01)*, pages 97–104. IEEE Computer Society Press, 2001.

[19] R. Kosara, S. Miksch, and H. Hauser. Focus+Context taken literally. *IEEE Computer Graphics and Applications*, 22(1):22–29, 2002. doi:10.1109/38.974515.

[20] M. Kreuseler, T. Nocke, and H. Schumann. A history mechanism for visual data mining. In *Proceedings of the IEEE Symposium on Information Visualization (InfoVis '04)*, pages 49–56. IEEE Computer Society Press, 2004.

[21] M. Krzywinski, J. Schein, I. Birol, J. Connors, R. Gascoyne, D. Horsman, S. J. Jones, and M. A. Marra. Circos: An information aesthetic for comparative genomics. *Genome Research*, 19(9):1639–1645, 2009.

[22] A. Lex, H.-J. Schulz, M. Streit, C. Partl, and D. Schmalstieg. VisBricks: multiform visualization of large, inhomogeneous data. *IEEE Transactions on Visualization and Computer Graphics (InfoVis '11)*, 17(12):2291–2300, 2011.

[23] A. Lex, M. Streit, E. Kruijff, and D. Schmalstieg. Caleydo: Design and evaluation of a visual analysis framework for gene expression data in its biological context. In *Proceeding of the IEEE Symposium on Pacific Visualization (PacificVis '10)*, pages 57–64, 2010.

[24] A. Lex, M. Streit, C. Partl, K. Kashofer, and D. Schmalstieg. Comparative analysis of multidimensional, quantitative data. *IEEE Transactions on Visualization and Computer Graphics (InfoVis '10)*, 16(6):1027–1035, 2010.

[25] A. Lex, M. Streit, H.-J. Schulz, C. Partl, D. Schmalstieg, P. J. Park, and N. Gehlenborg. StratomeX: enabling visualization-driven cancer subtype analysis. In *Poster Proceedings of the IEEE Symposium on Biological Data Visualization (BioVis'12) (to appear)*, 2012.

[26] M. McGuffin and I. Jurisica. Interaction techniques for selecting and manipulating subgraphs in network visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):937 –944, 2009.

[27] M. Meyer, T. Munzner, and H. Pfister. MizBee: a multiscale synteny browser. *IEEE Transactions on Visualization and Computer Graphics (InfoVis '09)*, 15(6):897–904, 2009.

[28] P. Neumann, S. Schlechtweg, and S. Carpendale. ArcTrees: visualizing relations in hierarchical data. In K. W. Brodlie, D. J. Duke, and K. I. Joy, editors, *In Proceedings of the IEEE Symposium on Data Visualization*, page 5360, Aire-la-Ville, Switzerland, 2005.

[29] C. North and B. Shneiderman. Snap-together visualization: A user interface for coordinating visualizations via relational schemata. In *Proceedings of the ACM Conference on Advanced Visual Interfaces (AVI '00)*, pages 128–135. ACM, 2000.

[30] S. Palmer and I. Rock. Rethinking perceptual organization: the role of uniform connectedness. *Psychonomic Bulletin and Review*, 1(1):29–55, 1994.

[31] S. E. Palmer. Common region: A new principle of perceptual grouping. *Cognitive Psychology*, 24(3):436–447, July 1992.

[32] F. V. Paulovich, F. M. B. Toledo, G. P. Telles, R. Minghim, and L. G. Nonato. Semantic wordification of document collections. *Computer Graphics Forum*, 31(3):1145–1153, 2012.

[33] H. Piringer, C. Tominski, P. Muigg, and W. Berger. A multithreading architecture to support interactive visual exploration. *IEEE Transactions on Visualization and Computer Graphics (InfoVis '09)*, 15(6):1113–1120, Dec. 2009.

[34] R. Rao and S. K. Card. The table lens: merging graphical and symbolic representations in an interactive focus + context visualization for tabular information. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '94)*, pages 318–322. ACM Press, 1994.

[35] N. Riche and T. Dwyer. Untangling euler diagrams. *IEEE Transactions on Visualization and Computer Graphics (InfoVis '10)*, 16(6):1090–1099, 2010.

[36] H.-J. Schulz. Treevis.net: A tree visualization reference. *IEEE Computer Graphics and Applications*, 31(6):11–15, 2011.

[37] H.-J. Schulz, S. Hadlak, and H. Schumann. The design space of implicit hierarchy visualization: A survey. *IEEE Transactions on Visualization and Computer Graphics*, 17(4):393–411, 2011.

[38] H.-J. Schulz, M. John, A. Unger, and H. Schumann. Visual analysis of bipartite biological networks. In *Proceedings of the Eurographics Workshop on Visual Computing for Biomedicine (VCBM '08)*, pages 135–142, 2008.

[39] Y. B. Shrinivasan and J. J. v. Wijk. Supporting the analytical reasoning process in information visualization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*, pages 1237–1246. ACM Press, 2008.

[40] M. Steinberger, M. Waldner, M. Streit, A. Lex, and D. Schmalstieg. Context-preserving visual links. *IEEE Transactions on Visualization and Computer Graphics (InfoVis '11)*, 17(12):2249–2258, Dec. 2011. Best paper award.

[41] 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.

[42] M. Streit, H.-J. Schulz, A. Lex, D. Schmalstieg, and H. Schumann. Model-driven design for the visual analysis of heterogeneous data. *IEEE Transactions on Visualization and Computer Graphics*, 18(6):998–1010, 2012.

[43] W. S. Torgerson. Multidimensional scaling: I. theory and method. *Psychometrika*, 17(4):401–419, Dec. 1952.

[44] A. M. Treisman and G. Gelade. A feature-integration theory of attention. *Cognitive Psychology*, 12(1):97–136, 1980.

[45] M. Waldner, W. Puff, A. Lex, M. Streit, and D. Schmalstieg. Visual links across applications. In *Proceedings of the Conference on Graphics Interface (GI '10)*, pages 129–136. Canadian Human-Computer Communications Society, 2010. Best student paper award.

[46] M. Waldner and D. Schmalstieg. Collaborative information linking: Bridging knowledge gaps between users by linking across applications. In *Proceeding of the IEEE Symposium on Pacific Visualization (PacificVis '11)*, pages 115–122. IEEE Computer Society Press, 2011.

[47] M. Wertheimer. Untersuchungen zur Lehre von der Gestalt. II. *Psychologische Forschung*, 4(1):301–350, 1923.

[48] W. Willett, J. Heer, and M. Agrawala. Scented widgets: Improving navigation cues with embedded visualizations. *IEEE Transactions on Visualization and Computer Graphics (InfoVis '07)*, 13(6):1129–1136, 2007.

[49] J. M. Wolfe. What can million trials tell us about visual search? *Psy-*

chological Science, 9(1):33–39, 1998.

[50] J. Wood and J. Dykes. Spatially ordered treemaps. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1348–1355, Nov. 2008.

[51] S. Zhai, J. Wright, T. Selker, and S.-A. Kelin. Graphical means of directing user's attention in the visual interface. In *Proceedings of the Conference on Human-Computer Interaction (INTERACT '97)*, pages 59–66. Chapman & Hall, 1997.

[52] C. Ziemkiewicz and R. Kosara. Laws of attraction: From perceptual forces to conceptual similarity. *IEEE Transactions on Visualization and Computer Graphics (InfoVis '10)*, 16(6):1009–1016, 2010.