Word-Sized Graphics for Scientific Texts

Fabian Beck, *Member, IEEE Computer Society,* and Daniel Weiskopf, *Member, IEEE Computer Society*

Abstract—Generating visualizations at the size of a word creates dense information representations often called *sparklines*. The integration of word-sized graphics into text could avoid additional cognitive load caused by splitting the readers' attention between figures and text. In scientific publications, these graphics make statements easier to understand and verify because additional quantitative information is available where needed. In this work, we perform a literature review to find out how researchers have already applied such word-sized representations. Illustrating the versatility of the approach, we leverage these representations for reporting empirical and bibliographic data in three application examples. For interactive Web-based publications, we explore levels of interactivity and discuss interaction patterns to link visualization and text. We finally call the visualization community to be a pioneer in exploring new visualization-enriched and interactive publication formats.

Index Terms—Sparklines, word-sized graphics, literature survey, text and visualization, interactive documents, scientific publishing.

1 Introduction

We, as researchers, publish papers like decades ago—a static text document enriched with figures and tables—although Web-based text documents would allow us to embed videos, 3D graphics, and interactive visualizations. Such media could augment the text with illustrating examples, supportive information, or quantitative evidence. However, these combinations also require the readers to jump between the text and the additional artifact. Psychology researchers investigate this *split-attention effect* [1] as part of *cognitive load theory* [2]: splitting the readers' attention could increase their working memory load and thus reduce their resources to process the actual content [3]. Hence, integrating the information into a single representation could avoid the split-attention effect and reduce cognitive load [1], [4].

Tufte [5] suggested *sparklines*—word-sized graphics—as an approach that embeds visualizations into text and creates a combined representation. For example, we can easily show a data-rich stock chart as part of this text 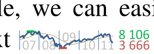

; the reader does not have to switch to a separate figure to read the information. The usage of word-sized graphics has become popular within spreadsheets and tables, significantly boosted by the integration of sparklines into Microsoft Excel and other software products. However, their embedding into text documents, especially scientific publications, falls behind and has not yet been sufficiently explored from a scientific perspective.

This work aims at studying the potential of word-sized graphics to enrich scientific texts. It both explores the state of the art and demonstrates new solutions. Figure 1 illustrates the envisioned integration process of text and word-sized graphics and previews application examples. In particular, our main contributions are:

- We perform a **literature review** with 140 publications to evaluate how word-sized visualizations have already been applied to scientific communication (Section 3).
- We discuss three generalizable **application examples** (Figure 1) that leverage the embedding of word-sized visualizations into the text of a publication (Section 4).
- We investigate levels of **interaction** for embedding word-sized visualizations into interactive text documents and propose reusable interaction patterns (Section 5).

We conclude our work with calling the visualization community to action (Section 6): *Explore visualization-enriched publication formats! Establish standards and best practices that support other communities!* We consider our work as a step in this direction.

2 Word-sized Graphics

Tufte [5] defined the term *sparklines* as "data-intense, design-simple, word-sized graphics." Although this definition is quite broad, the term *sparklines* is often understood in the narrow sense of *word-sized line charts* or *word-sized bar charts*, the latter sometimes also referred to as *sparkbars*. In this work, however, we consider the broader definition to cover all kinds of word-sized data visualizations: sparklines could be any data-intense visual representations at the size of a word. In particular, we use the term *word-sized graphics* to even include the coding of information using icon images. An alternative, very similar term is *word-scale graphics/visualizations*, defined by Goffin et al. [6] somewhat broader, for instance, allowing graphics that span paragraphs. Their definition also includes *emojis* [7]—emoticons (e.g., "☺") or pictures as placeholder for words (e.g., "☕ now?")—, which are wide-spread in contemporary written colloquial language. But since we focus on formal scientific communication only, we consider these as out of scope for the current work. Goffin et al. [6] discern between *data-driven* and *non-data-driven* graphics—our examples are all data-driven, even icon representations encode defined categorial attributes. Another related term is *glyph*, in visualization literature often understood as a "small visual object that depicts attributes of a data record" [8]. On the one hand, many word-sized graphics or sparklines can be considered as a special type of glyph, but the term *glyph* is broader, not limited to size or aspect ratio of a word. On the other hand, visualizations like small bar charts might not be understood as a glyph because they do not constitute a single visual object, but rather a collection.

Word-sized graphics could be embedded into different media: text, tables, user interfaces, source code, or even other visualizations. Since scientific publications primarily use text to

- *Fabian Beck is with University of Duisburg-Essen, Germany E-mail: fabian.beck@wiwinf.uni-due.de.*
- *Daniel Weiskopf is with VISUS, University of Stuttgart, Germany E-mail: weiskopf@visus.uni-stuttgart.de.*

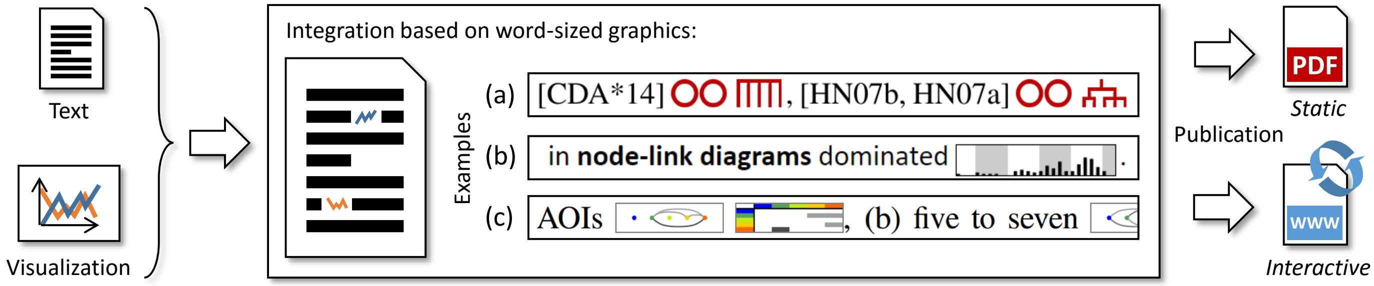

Fig. 1. Integration of text and visualization within scientific publications using word-sized graphics. Three generalizable examples demonstrate the practical applicability and potential of the approach: (a) classifying references with icon graphics, (b) discussing bibliographic data supported by interactive sparklines, and (c) summarizing eye movement data with word-sized visualizations.

communicate information, our main focus is on the integration of word-sized graphics into text. We consider them not as part of the language, but as an augmentation (much like a reference or a foot-note). Also, their usage in tables and figures is relevant in this work as a secondary means of communication within publications. In interactive Web-based documents, borders between tables, figures, and interactive visualizations blur. Even the graphics integrated into text could provide interactions, for instance, transforming them into a larger figure on click. We also discuss these interactive examples of scientific communication.

3 Literature Review

Despite being described in textbooks and implemented in software libraries, word-sized graphics are not necessarily applied by researchers. Since we want to study the usage of theses visualizations for scientific communication within texts, our starting point is to evaluate the popularity and usage scenarios in this kind of application. We aim at answering the following research question: *How are word-sized graphics used in scientific publications?* Hence, we systematically search for examples of word-sized graphics in the research literature and classify their form of usage. Relevant instances could be spread across multiple disciplines—the search should not be limited to visualization or computer science research.

We are not aware of any other survey on the usage of word-sized graphics in publications across scientific communities. Tufte [5] discusses general design considerations for sparklines. Others explore the interplay of sparklines and Gestalt laws [9], placement options for word-scale visualizations into text [10], [11], and the function of word-scale visualizations in documents [6]. A survey on the state of the art in glyph design is also related [8].

Please note that we ourselves apply word-sized graphics in the following to describe and quantify the collected literature. We also add word-sized graphics to give examples of representations used in the collected literature; these examples are representatives of a group of visualizations or replica of the original diagrams slightly adopted for consistency.

3.1 Methodology and Categorization Scheme

We performed a Google Scholar search with search terms *sparkline* and *Tufte* (retrieved: September 5, 2016). We used Google Scholar as a search engine because it lists publications from all disciplines and is one of the most extensive collections of scientific literature. We used the search term *sparkline* because it is—more than *word-sized graphics*—the accepted and unique term. We added the term *Tufte* to filter out noise and reduce the collection to those publications that pay credit to Tufte.

The search produced a result set of 542 publications. We excluded $^{38}/_{542}$ duplicates, $^{49}/_{542}$ non-English publications, and $^{93}/_{542}$ non-scientific publications (i.e., publications that were not published in a scientific context, for instance, user manuals, patents, talks, or blog posts). For the remaining $^{362}/_{542}$ publications, we tried to retrieve a full document copy (e.g., PDF or HTML version) and were successful in $^{316}/_{362}$ cases. In about half of these publications ($^{176}/_{316}$), we did not detect any original use of word-sized graphics, but just a reference to the term *sparkline*, for instance, as part of a literature overview; we also did not consider it as original use if sparkline-like graphics were considerably larger than the height of a few lines or carried axes and labels. We document these exclusion decisions in a table that is part of the supplemental material. The remaining $^{140}/_{542}$ publications are the final set on which the following classification is based. The literature collection together with the classification is available in the interactive literature browser SurVis [12]:

http://sparklines-literature.fbeck.com

We visually scanned each document for word-sized graphics. If the document was too long (e.g., PhD theses or books), we searched for the term *sparkline* within in the document. Based on the detected word-sized graphics, we classified the publication into the following main usage types (Category *type*):

- **Visualization Technique**: The work introduces a visualization technique using word-sized graphics.
- **User Study**: The work evaluates an approach based on word-sized graphics or their general usage in a user study.
- **Meta**: The work broadly discusses considerations for using word-sized graphics and draws general conclusions.
- **Scientific Communication**: The work uses word-sized graphics only to communicate data—the visualization itself is no main contribution.

The last type is the one we are most interested in within the scope of this work. While we briefly review publications of the other types as well, we focus the following literature analysis onto the examples assigned to the *communication* type. Our paper—both making general observations and using word-size graphics to communicate results—would be classified as *meta* and *communication*.

We further categorize publications according to the scientific domain they target (Category *domain*); this domain does not refer to the scientific community in which the work was published, but

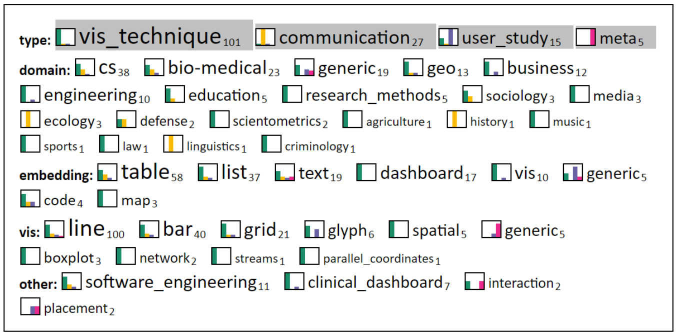

Fig. 2. Word cloud summary of keywords assigned to publications; subscript numbers and font sizes encode frequencies; usage type keywords are assigned a unique color: *visualization technique* ■, *communication* ■, *user study* ■, and *meta* ■; word-sized graphics next to every keyword show the overlap of the respective keyword with the usage types (e.g., for *domain:cs* , the yellow bar indicates that about a third of these are *communication* publications); created with SurVis [12].

the domain of the data visualized in the word-sized graphics. We also discern different forms of embedding for word-sized graphics (Category *embedding*), for instance, tables, lists, texts, and other representations. Finally, we classify the visual encoding used in the word-sized graphics (Category *vis*), for example, line charts, bar charts, or spatial encodings.

The classification was the result of an iterative coding process, refining categories during classifying new publications and reclassifying previous publications when adapting the overall scheme. We assigned at least one keyword for the categories *type*, *domain*, *embedding*, and *vis* to every publication in our collection. Only in case a publication could not unambiguously be associated with a single keyword, we assigned multiple keywords per category. To highlight aspects beyond this classification scheme, some publications carry additional independent keywords.

3.2 Results

The collected publications per year $^{21}_{0}$ (plotted from 2006 to 2016) are a general indicator for the popularity of word-sized graphics within the scientific community: After Tufte [5] coined the term in 2006, sparklines had a direct impact as they are referenced immediately (on average, 9.0 publications per year in 2007 to 2009). Over the years, their popularity varied somewhat but steadily increased (on average, 20.0 publications per year in 2013 to 2015 ; 2016 is not considered here because the data might not yet be complete).

Figure 2 quantifies the keywords assigned to the publications as a word cloud, structured by keyword categories and encoding keyword frequencies in subscript numbers and font sizes. Type keywords are highlighted and word-sized graphics encode the relationship of each keyword to these main discerning types. In the following, we discuss this classification in detail and use the keyword categories to structure the analysis. At the end of each analysis part, we formulate an *observation* that condenses what we believe is the central outcome of the respective analysis. Please note that these statements only reflect our interpretation of the data, but might not be the only valid conclusion that can be drawn.

3.2.1 Usage Type

Like indicated in Figure 2, a majority of $^{101}/_{140}$ publications in our literature collection uses word-sized graphics as part of a *visualization technique*. Some publications only marginally discuss

TABLE 1
Existing examples where word-sized graphics are used in publications for scientific communication.

	Visualization	*Short Description*	*Domain*	*Embedding*
[13]		Query event sequence	Comp. Sc.	**Text**
[14]		Discussion subjects	Comp. Sc.	List
[15]		Participant rating	Comp. Sc.	Table
[16]		Cost across time	Defense	**Text** + table
[17]		Retention/completion rates	Education	Table
[18]		Levels of suspicious activity	Comp. Sc.	List
[19]		Infection rates	Bio-medical	Table
[20]		Water chemistry species	Ecology	Table
[21]		Transaction costs	Ecology	Table
[22]		Volatilization of hydrocarbon	Bio-medical	Table
[23]		Weights of signaling metrics	Bio-medical	Table
[24]		Phasic pressure traces	Bio-medical	Table
[25]		Criticality distributions	Comp. Sc.	Table
[26]		IO rates	Comp. Sc.	Table
[27]		Score by amount of smoking	Bio-medical	List
[28]	89% 63%	Population statistics by age	Sociology	Table
[29]		Concentration of positives	Comp. Sc.	Table
[30]		Temperature signatures	Bio-medical	List
[31]		Topic frequency in publications	History	Table
[32]		Topic importance trend	Comp. Sc.	Table
[33]		Litter fall	Ecology	Table
[34]		Term likelihood per category	Linguistics	**Text**
[35]		Web service characteristics	Comp. Sc.	Table
[36]		Answer distributions	Bio-medical	Table
[37]		Attraction schemas	Bio-medical	List
[38]		Participant agreement	Comp. Sc.	Table
[39]		Observed behavior per user	Comp. Sc.	**Text** + table

the integration of word-sized graphics. In contrast, such graphics are in the focus of other publications, for instance, to encode social networks in a matrix [40], debug electronic circuits with an interactive tabular representation [41], summarize questionnaire responses in small multiples [42], show the evolution of keywords in a tag cloud [43], compare time series in a matrix [44], or visualize geo-located attributes over time [45]. Further, $^{15}/_{140}$ publications evaluate word-sized graphics in *user studies*, for instance, the combination of quantitative and qualitative data in time series visualizations [46], the effectiveness of a medical data display based on word-sized graphics [47], the influence of graphics embedding into text onto reading behavior [10], or the performance of word-size data representations within tables [48]. We classified $^{5}/_{140}$ publications as *meta*, discussing the placement of word-sized graphics in text [11], user interactions with word-sized graphics [49], aspect ratios of diagrams and sparklines [50], and *micro visualizations* as different means of visualization to augment texts [51], [52].

In contrast to the rather wide use as part of visualization approaches, word-sized graphics are less common for communicating scientific data and results ($^{27}/_{140}$ publications). Table 1 lists these instances in detail, also providing a graphical classification of each. While the short description of the visualizations indicates a great variety of approaches, the graphics have in common that they only play a subordinate role in the publication. Unlike in the publications classified as *visualization technique*, here, word-sized graphics are just used to communicate some information, but do not belong themselves to the core contributions of the work. The following analysis parts look into more details of these examples.

Observation: *While word-sized graphics are regularly applied within visualization approaches and systems, they are yet rarely used for communicating data within scientific publications.*

3.2.2 Application Domains

The application domains for word-sized graphics are diverse. As shown in Figure 2, our literature collection covers areas such as computer science, bio-medical applications, geography, business, engineering, and more. Among these, computer science and bio-medical applications are most prominent ($^{38}/_{140}$ and $^{23}/_{140}$ publications). While the variety of topics is quite high also within these disciplines, we observed some clusters of publications: for computer science, related to software engineering topics ($^{11}/_{38}$ publications), and for bio-medical publications, related to clinical data dashboards showing patient records ($^{7}/_{23}$ publications). Also, notable numbers of publications can be found in geographic applications ($^{13}/_{140}$), business ($^{12}/_{140}$), and engineering ($^{10}/_{140}$). We marked $^{19}/_{140}$ publications as *generic* regarding their application domain because they cover visualization research applicable to many domains, for instance, *user study* or *meta* publications.

For the publications classified as *scientific communication*, Table 1 shows a similar distribution of application domains with many examples in computer science and bio-medical applications. The short descriptions provide more context and illustrate the high variance of represented data. We could only detect two clusters of publications using word-sized graphics for a similar purpose: one for summarizing the distribution of participants' answers [15], [36], [38], the other representing topic frequencies or importance [31], [32]. From an abstracted perspective, the visualized data in all examples only has in common that it was measured in some sort of experiment.

Observation: *The broad coverage of application domains shows that word-sized graphics have the potential to be used within publications in all (data-driven) sciences.*

3.2.3 Embedding

Word-sized graphics, due to their small size, can be embedded into different artifacts. Among the reported use cases in our collections, their usage as small visualizations in cells of the table dominates ($^{58}/_{140}$ publications, cf. Figure 2). Showing them as small multiples in a list—similar to a table with one column or row—is a variant ($^{37}/_{140}$). The integration of word-sized graphics into text, although already proposed by Tufte [5], is less common, but at least $^{19}/_{140}$ publications follow this approach. While these three basic embeddings work with static and interactive representations alike, there also exist a number of embeddings that can be applied specifically within interactive user interfaces: using word-sized graphics to build a dashboard ($^{17}/_{140}$), augment a larger visualization ($^{10}/_{140}$), enrich source code in an editor ($^{4}/_{140}$), or add extra information to a map ($^{3}/_{140}$).

Also among the examples of scientific communication, embeddings in tables and lists are prevalent ($^{25}/_{27}$ examples, cf. Table 1). As highlighted in bold font in the table, $^{4}/_{27}$ publications, however, use an embedding into text ($^{2}/_{27}$ in co-occurrence with *table*): Adar et al. [13] show temporal event sequences of search queries integrated into the text. Boehmke [16] visualizes evolving costs of different categories within text and tables. Potts [34] shows term likelihood per rating category of movie reviews. Ying and Robilliard [39] summarize specific behavior of users to discuss typical user strategies. Comparing the three embeddings (table, list, and text), we observe that the ratio of communication examples (height of bar) is larger among table representations than for lists and texts; a tabular embedding might be most straightforward to use in a paper.

Within publications not classified as usage type *communication*, further relevant examples of text embedding exist: Brandes and Nick [40] use word-sized graphics to represent the evolving relationships of pairs of individuals. For song lyrics, Oh [53] embeds visualizations above each line that encodes melody and beat. Tinkelman [54] introduces a word-sized graphic designed to encode losses and profits of companies and demonstrates this application within four written statements. From a general visualization research perspective, Goffin et al. [11] explore the options for placing word-sized graphics into text and identify seven cases ranging from in-line placement as practiced in this work to inter-line placement and overlays. Goffin et al. [10] further investigate in a user study how these placement options interact with reading behavior. In most other cases that we classified as *text embedding*, word-sized graphics are used in the text in only few instances or the text embedding just plays a marginal role.

Observation: *Word-sized graphics included into tables and lists are much more common than examples integrated into text, both in publications describing a visualization technique and those using them for scientific communication.*

3.2.4 Visualization

Finally, we compare the visual encoding used in the word-sized graphics. Figure 2 clearly shows that most examples relate to simple line or bar charts ($^{100}/_{140}$ and $^{40}/_{140}$ examples, with an overlap of $^{17}/_{140}$ examples that show both; please note that we also consider area charts as variants of line charts). Similar to the original examples provided by Tufte [5], some of them are enriched with additional markers. Another simple encoding employs grids to encode, for instance, a sequence of states in colored cells in one or multiple rows ($^{21}/_{140}$). However, there also exist approaches demonstrating that word-sized graphics are not limited to these simple diagram types, for instance, graphics that encode spatial trajectories or densities [55], [56], [57], stacked quantities that form streams [58], small representations of boxplots to display statistical distributions [59], [60], [61], or glyphs that encode multivariate properties [9], [40]. Even parallel coordinates can be represented [42], and networks in simplified node-link representations [55], [62] or adjacency matrices [55]. The publication years of the respective publications show that these alternative representations are rather new, mostly suggested in the past four years.

In the examples classified as *scientific communication*, also line and bar charts dominate (cf. Table 1). In addition to standard lines, some lines carry markers to indicate specific points and periods in time or average, minimum, maximum, or last value; one example plots multiple lines in a single chart [28]. Bar charts occur with monochrome bars and bars colored by category or value; bars can be also tailored and carry markers [23]. Two of the communication examples also use a grid-like visualization to show temporal sequences of predefined topics [14] or encode ternary vectors [37]. Abstracting from individual representations, a commonality among most examples is the encoding of time or sequence on the horizontal axis and of a measured value on the vertical axis.

Observation: *Line and bar charts are frequently applied in word-sized graphics, using time or sequence as horizontal and value as vertical dimension. It is also possible to use other types of visualizations as word-sized graphics, but this is not yet much explored for scientific communication.*

3.3 Discussion

The literature review provides an impression of the current use of word-sized graphics in scientific publications. Word-sized graphics are popular in general, but their application for scientific communication is still quite infrequent. However, we may not have found all publications, in particular, if the authors were not aware that they are using a form of visualization referred to as *sparklines* or do not acknowledge Edward Tufte. Whereas our classification scheme is simple and does not require much interpretation, still in some of the retrieved publications, the classification was ambiguous. Others might have categorized individual publications differently or would have used another classification scheme, potentially resulting in different outcomes of the analysis. As discussed above, the communication examples found cover mostly usage in tables and simple line and bar charts. Word-sized graphics are not much used in practice within the text of scientific publications to communicate data and results. Reasons for this gap can be manifold, for instance, (i) word-sized graphics could be difficult to integrate, (ii) authors were afraid that reviewers or readers are too conservative regarding their layout preferences of a scientific publication, or (iii) data is too complex to be represented as word-sized graphics. In contrast, we believe that the integration is both feasible and promising in many applications.

4 Application Examples

To further explore the usage of word-sized graphics for scientific communication, we study a number of examples that demonstrate positive use cases. Since the usage of sparklines in tables and lists is comparably well established and quite similar to normal tables and lists, we focus here on examples where visualizations are directly embedded into text. For text, a good integration is harder to achieve but, if successful, promises to avoid a split-attention effect [1].

The three examples we discuss in the following cover a variety of graphics, from icons for simple classification to visualizations of rather complex data. Unlike the state of the art applied for scientific communication (cf. Table 1), these examples demonstrate that word-sized graphics are not limited to line and bar charts. Their areas of application relate to general topics, like communication of bibliographic data and experiment results, and could be reused in a similar way in many domains and publications. They illustrate one possible solution tested in practical application, but do not claim to be the best solutions. Whereas we motivate design considerations and suggestions, they are not backed by empirical evidence. Extensive user studies would be necessary to come up with guidelines and more detailed recommendations. Parts of the discussed examples have already been introduced elsewhere [55], [63], [64], however, without particularly discussing their use within texts of scientific publications. For each example, we first introduce the visualization, then generalize it, and finally discuss advantages and limitations.

4.1 Classification of References

When discussing related work, authors need to highly condense information about the referenced publications in one sentence or a few. Often, it is hard for the reader to follow the discussion and understand the main differences between the briefly discussed approaches. This problem is particularly critical for state-of-the-art reports or other literature surveys where large parts of the paper consist of content like this. Typically, the authors of such a publication have already carefully structured and classified the referenced approaches to give an overview and point out differences. While the classification of publications might be communicated through page-filling tables, it requires some searching (and maybe page flipping) to match the reference in the text with the ones in the table. Hence, using icons for the classification and embedding them into the text directly next to the references provides an alternative or at least additional encoding. Although debatable whether these icons should be called *word-sized graphics* or *sparklines*, we consider this example because these icons are a visual encoding of data—an assignment of publications to categories.

4.1.1 Example

In a survey on visualizing group structures, Vehlow et al. [64] tested this approach. They first introduced and explained the icons, here, the ones referring to the type of group structure represented:

> "The references are [..] marked with the respective icons: flat , or hierarchical , disjoint or overlapping ."

Within the discussion of the surveyed publications, these icons are used to classify the publications, often providing an additional information that is not described explicitly in the text, for instance:

> "Group nodes are connected by visual links if any of their members are related [CDA*14] , [HN07b, HN07a] ."

Further icons are applied, such as, for the type of overlap , the graph visualization paradigm , the encoding of time , and the usage of color coding . One of the tables in the document acts as a legend for the icons in addition to their introduction in the text. Icons are only added if the classification is helpful as an additional information in the specific context. Within table headers and figure captions, icons are used as well to specify listed approaches.

4.1.2 Generalization

This approach is easy to generalize even if it was designed for a specific example. In essence, each reference is assigned a number of categories, each expressed with a different icon. While it is an important goal to make these icons as self-explaining as possible, like for acronyms, one should always add a textual explanation before using the icons for the first time. In context of visualization literature, finding self-explaining icons might be easier than in other domains because one can often use stylized small versions of basic visualization approaches; for instance, Kucher and Kerren [65] provide another example of such icons for text visualization. However, other domains have their visual languages and accepted metaphors that can be used as a basis for designing the icons. Often, visual identifiers could be borrowed from icons used in software systems popular in the respective domain. While it might be possible to use arbitrary icons that are not obviously related to the category they represent, we do not expect that readers would be willing to learn this artificial visual language—the overhead would more than compensate the gain.

4.1.3 Discussion

Classifying references with icons is a lightweight way of using word-sized graphics within scientific publications. The encoding is easy to understand and introduce. The icons do not require much space or an enlarged version in the paper. The icons only occur when references, as non-language artifacts, enrich the text with extra information anyways—the reader is already accustomed to intermission of reading in this context. As main advantages we see that the icons (i) show additional illustrations of the referenced publications, (ii) provide indicators of what are similarities and differences between referenced publications, and (iii) allow the reader to quickly search in the document for specific aspects represented by icons. We recommend focusing on the most important aspects and leaving out icons if the classification is obvious or redundant in order to avoid visually overloading the paper.

4.2 Bibliographic Analysis

While the approach discussed above classifies individual references qualitatively, we could also consider a more quantitative approach of discussing and summarizing scientific literature—counting structured publications in a literature collection and investigating temporal trends.

4.2.1 Example

You have already seen examples of this use case as part of our literature search (Section 3). In particular, we use three different encodings that aggregate publications to quantities:

- **Relative Publication Numbers** (e.g., 42/140): Two overlapping bars show the number of publications in relation to a reference number (i.e., all publications of the collection or of a previously defined subset). We also provide the absolute numbers written as a quotient.
- **Publication Timelines** (e.g., 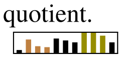$^{21}_{0}$): Publication frequencies per year plotted as bar charts give insights into impact and trends. At first occurrence, we augment this representation with labels for minimum and maximum value. Coloring the bars helps us highlight different periods or subsets of the publications. These timelines are similar to the line-based visualization of topic frequency used by Milligan [31].
- **Keyword Overlap** (e.g., ): Showing the usage types of a certain publication set identified by a keyword, we illustrate the underlying subset structure of these publications. These graphics were created with the SurVis system [12]. We now show that these can be used also to enrich a textual description of a literature collection.

4.2.2 Generalization

Since the applied visualizations are simple, they obviously generalize to reflect arbitrary relative quantities, aggregated number of events or frequencies on a timeline, or overlap of item sets. One could discuss other document or media collections in a similar style, for instance, patents, historic texts, works of art, or music. Another relevant extensions could be to include more metadata into the discussion, such as information about authors, publishers, or citations. Whereas some of this information can be represented with the proposed visualizations already (e.g., citations per year $^{133}_{0}$, 2006–2016 of the book *Beautiful Evidence* by Tufte [5]; retrieved from Google Scholar, September 30, 2016), other data requires new word-sized visualizations (e.g., citation or co-author networks).

4.2.3 Discussion

In analogy to a reference or a parenthetical explanation, we add the visualizations usually at the end of the respective sentence or clause it refers to. However, since relative publication quantities can be considered as a number, we also embed them within the sentence like a number. One might argue that our encoding of publication numbers is superfluous because a single number as plain text is sufficient. However, we assume that often the absolute number of publications within a category is important, as well as the percentage with respect to a total number. If just providing the absolute number, the reference is not always clear; if just providing a percentage value, the reader needs to calculate the absolute number. The graphical representation further provides the advantage to quickly search for high quantities and compare subsets by scanning through the quantities.

We pay attention that we do not change the reference of the graphics within a certain context, that is, the reference number of the publication quantities, the normalization and time period of the timelines, and the color coding of the overlap visualizations. Otherwise, the readers might misinterpret the graphics or each instance would need to be accompanied with lengthy explanation, partly destroying the embedding effect.

4.3 Eye Movement Visualizations

Eye tracking is a specialized, but increasingly used empirical evaluation methodology, for instance, in psychology, human-computer interaction, and visualization research. Reporting the results of such studies is difficult because we have to deal with spatio-temporal patterns of eye movements. While a variety of visualizations is available for this purpose already [66], an open research question was if these visualizations could also be transformed into word-sized representations and stay readable. Such a representation allows us to also embed eye movement visualizations into text, which we explore in this example. The variety and complexity of the visualizations show that the use of word-sized graphics is not limited to line and bar charts.

4.3.1 Example

In previous work, Beck et al. [55], [63] discuss how to transform existing eye movement visualizations into word-sized representations. While demonstrating the usage of the resulting word-sized visualizations in tabular interfaces, the authors just casually integrated the word-sized visualizations into the text of the publication for giving an example or performing a case study. In the following, we discuss this usage scenario in more detail, but first provide a brief introduction to a selection of the suggested visualizations.

Eye movement data consists of *fixations* and short transitions in between them, called *saccades*. Plotting the spatial position of these fixations connected with straight lines according to their sequence of occurrence creates a simple scan path visualization that can be shown at size of a word  (here, a color map 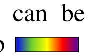 indicates the progress of time). When spatially aggregating the fixations and plotting their density, we obtain an attention map visualization . Putting time in the focus of the visualization by using the horizontal axis as a timeline, we could plot another metric to the vertical axis, for instance, one of the spatial dimensions of the fixations . To simplify the analysis and raise the level of abstraction, analysts identify *areas of interest* (AOIs) and transform the fixation sequence to a sequence of AOIs, which can be visualized on a similar timeline representation 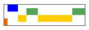 (here, colors and vertical position identify an AOI). Abstracting further, we transform the sequence of AOIs into a network of AOI transitions and visualize it as a graph, for instance, in a node-link representation 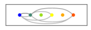 or an adjacency matrix 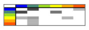.

Each of the presented visualization represents the eye movement for a single participant in a single task. Hence, integrating them into text is useful for illustrating an example of a typical fixation sequence or highlighting a specific outlier. In contrast to traditional eye tracking analysis that focuses on statistically aggregating the data, this integration of word-sized graphics much better supports a qualitative analysis for studying individual behavior in detail. Even contrasting a few individuals is possible, for instance, two participants that show a similar temporal sequence of AOIs versus another two participants forming a different group [63]:

> "Participants in Group (a) found a path from origin to destination quickly and did not verify it fully  . In Group (b), the participants performed a final verification more thoroughly 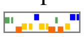 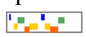."

If more than two or three instances are required to illustrate a finding, the word-sized graphics may consume a full line; they could be treated like a full-line formula and be centered, for instance, a set of visualizations showing different aspects for the same participant

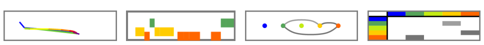.

This representation facilitates enlarging the graphics to some extent. Of course, the visualizations could also be embedded into a table. To show details, a few of the examples might be displayed in regular size enriched with labels as a figure.

4.3.2 Generalization

Although specifically developed for eye movement data, most of these visualization can be easily applied to other, related data. First, interaction data is very similar to eye movement data because interactions, like fixations, can be associated with a timestamp and screen location [55]. Moreover, eye movements are a special form of trajectories—hence, the visualizations can also be used to represent mobility patterns or other spatio-temporal data. Parts of the proposed visualizations show small networks, and hence, are applicable to any other kind of network, for instance, software dependencies or social networks.

4.3.3 Discussion

When representing data at the size of a word, in particular complex data, visual scalability and clutter quickly becomes an issue. For the suggested word-sized eye tracking visualizations, we observe problems like overplotting  or coarse data granularity 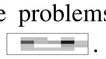. The AOI-based visualizations become cluttered when adding more AOIs. But a larger visualization does not always solve the scalability problem: In case of overplotting, the problem becomes only little better, if at all, when the graphic is enlarged by a magnitude, for instance, by a factor of nine regarding the area

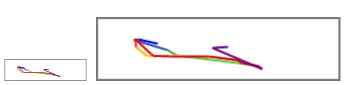.

Instead, we might divide the data into smaller chunks or use visualizations that abstract the data to a higher level of aggregation. For the spatial visualizations, another problem is that word-sized graphics usually have a wide landscape format, not necessarily similar to the aspect ratio of the stimulus. Also, the visualizations are too small to have the original stimulus shown in the background like possible in larger representations. Despite these issues, however, we were surprised ourselves how many regular eye tracking visualizations can be transformed to a representation at the size of a word without losing too much of the perceivable information and visual scalability. A more detailed discussion of visual scalability of the word-sized representations is available in previous work [55], as well an expert review of the different eye tracking visualizations [63].

5 Interaction

Like regular information visualizations, word-sized graphics become more powerful when the users or readers can interact with them—users might retrieve details on demand, explore relationships through brushing and linking, or adapt the visual encoding to their needs. There have only been few works yet that discuss interactions with word-sized graphics in detail: In context of debugging electronic circuits, Frishberg [41] describes the design challenges and user expectations for interactive sparklines embedded into a tabular interface. Goffin et al. [49] investigate how word-sized graphics embedded into text can be interactively transformed from a *document-centric view* (i.e., focusing on the text with word-sized embeddings) to a *visualization-centric view* (i.e., focusing on the visualization, potentially enlarged and interactive). They also describe intents, techniques, and scope of interactions on word-sized representations. In other examples, word-sized graphics are integrated into user interfaces, but interacting with these graphics is not discussed in detail.

Quite orthogonal to previous work, we discuss and classify levels of interactivity. We do not focus only on interaction regarding the visual representation but also the textual content and how these representations interact. This is particularly relevant for scientific communication because of the importance of the text in a publication. We first describe an interactive example that illustrates how interaction provides extra value within a text embedding of word-sized graphics. We later generalize interactions to patterns.

5.1 Example

We extend the timeline representation of our bibliographic analysis example (see Section 4.2) to an interactive version:

http://biblines-example.fbeck.com

We created a new textual description summarizing literature data publicly available in context of a survey on dynamic graph visualization [67]. The example combines text, word-sized timelines, and a regular-size timeline diagram. While Figure 3 shows a static overview of the example, Figure 4 illustrates the interactions:

- **Hovering a word-sized graphic** marks the clause it refers to and also shows the displayed data in the regular-size timeline as an overlay with blue bars (Figure 4 (a)).
- **Clicking on a word-sized graphic** makes the overlay persistent—the previous timeline data is replaced (Figure 3). The word-sized graphic currently shown enlarged is marked with a light blue background.

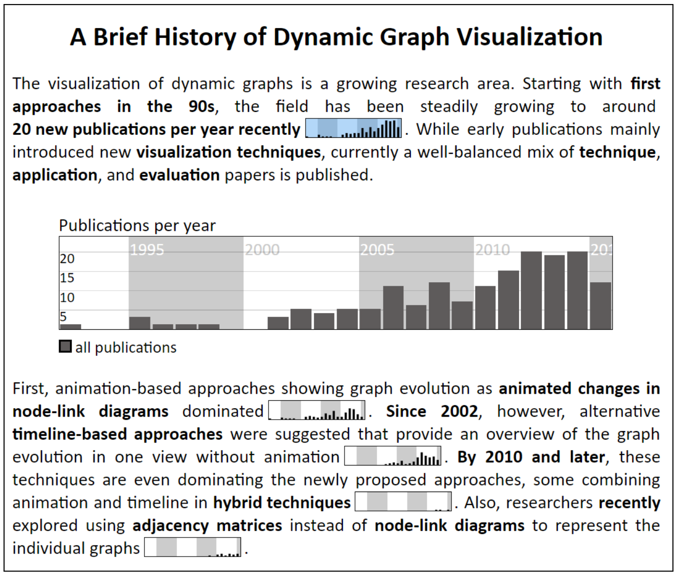

Fig. 3. Interactive Web-based integration of word-sized graphics into text discussing a bibliographic analysis of a literature collection on dynamic graph visualization [67].

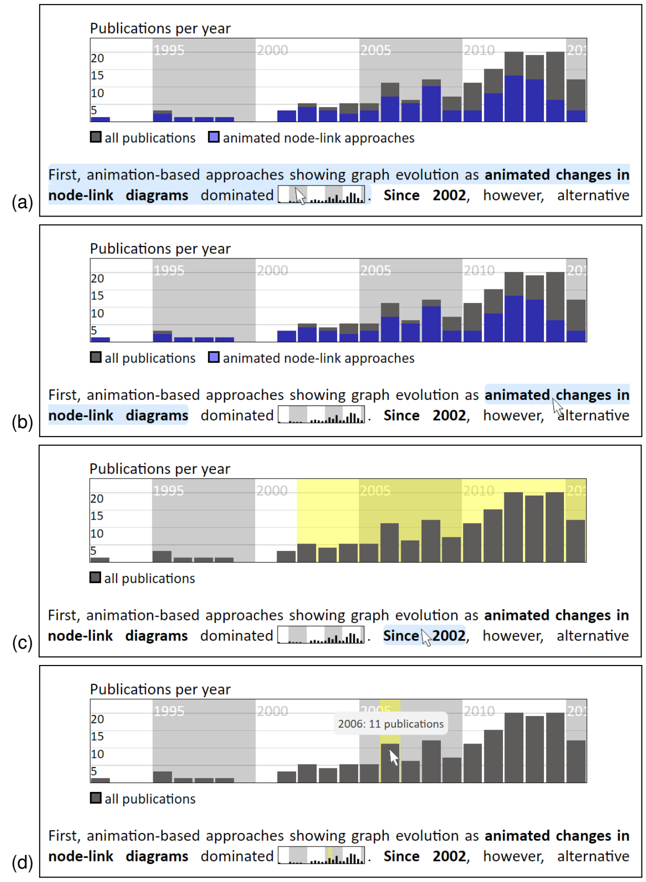

Fig. 4. Interactive interplay of text, word-sized graphics, and a diagram when hovering (a) a word-sized graphic, (b) highlighted text representing a subset of the data, (c) highlighted text representing a time period, or (d) a bar in the diagram.

- **Hovering a word in bold font** highlights a subset of the data. If the text refers to a time-independent subset of the data, this data is plotted as an overlay with blue bars on the enlarged diagram (Figure 4 (b)). If the text describes a time period, the respective period is highlighted with a yellow background in the enlarged diagram and in the word-sized graphic attached to the respective clause (Figure 4 (c)).
- **Hovering a bar in the diagram** provides details as a tooltip dialog and marks the respective year with a yellow background color, in the enlarged diagram as well as in all word-sized graphics (Figure 4 (d)).

Hence, after clicking on a word-sized graphic, users may explore the timelines in detail. When hovering a word-sized graphic or parts of the text, users can compare the temporal distributions of two publication sets. Hovering years or textual descriptions of temporal periods allows users to explore and compare temporal intervals across different representations. Since the word-sized graphics are explained by the regular-size diagram, the text is readable in static form already. However, the interactive version increases the value of the augmentation because it allows retrieving details and enables comparison. Hence, readers can use interactions, but do not have to do it to understand the text.

5.2 Levels of Interactivity

The design space of interactivity for word-sized visualizations spans, as demonstrated, from static publications with no interactions to highly interactive examples. This span is not only a difference in quantity of available interactions, but also covers different qualities and scope of interaction. In particular, we identified the following three levels:

- **No Interaction (Level 0)**: In their basic usage and as demonstrated above, word-sized graphics embedded into text do not need to be interactive at all. If not overly complex or condensed, the small-scale visualizations could already provide valuable extra information. However, since the graphics are usually too small to carry axis labels and captions, authors need to take care to provide sufficient explanations in the text surrounding the word-sized graphics.
- **Local Interaction (Level 1)**: A first step when adding interaction to word-sized graphics is providing details on demand, for instance, a tooltip dialog with explanations or showing an enlarged and labeled version of the graphic. Also, one could link the text and the graphics interactively; for example, when hovering a related text fragment or graphic, the respective other element gets highlighted. These simple interaction techniques share that their effect is local, only adding or highlighting information in the local environment of the graphic.
- **Global Interaction (Level 2)**: Going a step further, global interactions describe connections beyond the local scope of a graphic. In the above example, instances of a global interaction are a consistent linking across multiple word-sized graphics (cf. hovering a bar in the diagram) or a connection with an enlarged graphic (cf. hovering a word-sized graphic). Although arbitrary global connections are possible, we recommend keeping the scope somewhat restricted to not overwhelm the user with changes.

Hence, according to this classification, the examples as provided within the PDF version of the paper are obviously non-interactive (Level 0). In contrast, the interactive example described above (cf. Section 5.1) contains local and global interaction—thus, it is classified as Level 2. Transitions from document to

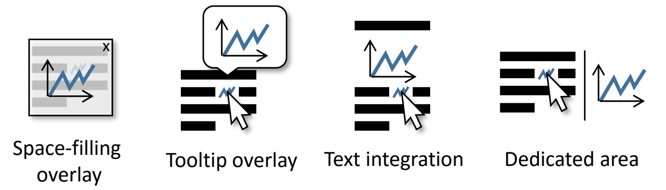

Fig. 5. Design space for showing an enlarged visualization as interactive details-on-demand of a word-sized graphic.

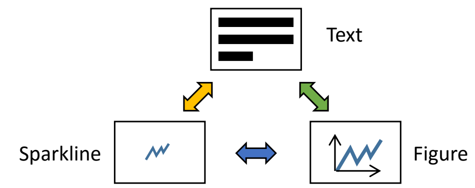

Fig. 6. Possible interactive links between text, sparklines, and figures implementing an *visualization–text interaction* approach.

visualization as described by Goffin et al. [49] applied to a single graphic are an example of local interactions (Level 1). A Level 2 interaction occurs when more than one graphic is involved in an interaction—discussed as *interaction scope* by Goffin et al. [49, Section 3.4].

5.3 Interaction Patterns

We instantiate the interaction levels as reusable patterns applicable for scientific communication. To this end, we sketch two patterns, one for each level. Please note that these patterns are only examples, other instantiations are possible. Nevertheless, we believe that these patterns are applicable and useful in many scenarios.

5.3.1 Details-on-Demand Interaction

A straightforward local interaction technique (Level 1) for improving the understandability of word-sized graphics is to provide details on demands. While some textual extra information showing labels or numbers would be a basic version of such details, these details could also contain an enlarged, possibly interactive version of the word-sized representation. For the placement of the details, there are the following options, which Figure 5 illustrates:

- **Space-Filling Overlay**: A screen-filling window or graphical layer containing the details overlays the text.
- **Tooltip Overlay**: A small tooltip dialog overlays part of the text next to the referenced element in a way that the referenced element does not get occluded.
- **Text Integration**: Details are blended in and displace part of the text—the text layout changes.
- **Dedicated Area**: Details are shown in a side bar or any other screen area that is specifically dedicated to this purpose and does not overlap the text.

If the details show an enlarged graphic, it should not differ much from the word-sized representation so that readers immediate see the congruence. We recommend keeping the aspect ratio consistent and just adding details such as axes and labels. In case users want to interact with this enlarged version, showing the details only when hovering is not sufficient—the users need to make the details persist on screen (e.g., using click) before they can interact. Implementing such interactions is possible with all placement options; only for tooltip overlays, one needs to prevent that they disappear when the mouse leaves the reference. For touch interfaces, hovering needs to be replaced by *tap* or *tap-and-hold*.

5.3.2 Visualization–Text Interaction

As an example of global interaction (Level 2), we abstract the interactive example described in Section 5.1 to a generic approach that we call *visualization–text interaction*. The focus of this approach is interactively linking visualizations and text. We assume—like in the example—that the interactive document does not only contain text and word-sized graphics (here, called *sparklines* for simplification) but also regular-size figures. Hence, interactive, bidirectional links could be integrated between all three kinds of representations like illustrated in Figure 6:

- **Sparkline–Text Interaction**: Linking the word-sized graphics and the text usually refers to a local interaction that improves the integration. For instance like in the example, when hovering on a word-sized representation, the text fragment it refers to becomes highlighted (⇒). The other way round (⇐), highlighting could be implemented in a similar way; in the example, we highlight a period in the related word-sized graphic when hovering a bold-font text fragment referring to a time period. Although it is possible to implement global sparkline–text interactions, we recommend keeping this link local because otherwise the close integration of text and word-sized representation would be partly destroyed.
- **Text–Figure Interaction**: Text and figure can be connected in a global interaction: a traditional reference of a figure in the text (e.g., “Figure 6”) could be made interactive by highlighting the figure when hovering the reference or the related text fragment (⇒), and vice versa, all text discussing the figure when hovering the figure (⇐). This is particularly useful if the document contains several figures. Since we use only one figure in our example, we just implemented a link referencing to different parts of the figure—when a description of a time period is hovered, the respective period gets highlighted in the figures.
- **Sparkline–Figure Interaction**: A set of interactive representations in word-sized graphics and figures can be considered as a multi-view visualization. Hence, typical multi-view interaction approaches like *brushing and linking* can be applied: Selecting something in a figure could highlight the respective visual elements in the word-sized graphics (⇒), in our example, a bar referring to a year. Regarding the opposite direction (⇐), however, selecting elements in a word-sized graphic is difficult due to their small size. Still, select operations might refer to all items represented in the graphic. In the example, hovering a word-sized graphic overlays the displayed subset of the data in the figure.

Through these links, text and visualizations become an integrated unit, mutually enriching each other. Like in a regular publication without word-sized graphics or interaction, the dominating usage strategy would still be linearly reading the text and occasionally using visualizations to better understand. However, we believe the interactive version has advantages on top of a better information integration: (i) references between representations become clearer and easier to trace, (ii) interactive links create a connected visualization showing different facets of the data, and (iii) interaction supports alternative reading strategies, such as studying figures first and then looking for the respective text.

6 Conclusion and Call to Action

This paper discussed the potential of embedding word-sized graphics within scientific texts. A survey of the state-of-the-art usage revealed that word-sized graphics are both flexible in what data they represent and universally applicable across scientific domains. However, authors have just started to use them for scientific communication. We demonstrated generalizable use cases leveraging word-sized graphics in a variety of applications. Interactive publications could further exploit these kinds of visualizations, providing details on demand and interactively linking text and visual content.

Call to Action: We believe that the visualization community can and should play a leading role in establishing word-sized graphics and other visualization-based enrichments for scientific texts. Visualization researchers have the expertise to provide positive examples and set standards. We call members of the community to action, in particular, (i) authors to experiment with better integration of text and visualizations in their papers and explore the use of interactive publication formats, (ii) reviewers to be open for new ideas and value experimental publication formats, and (iii) editors, conference chairs, and publishers to accept and technically support the submission of publications where borders between text and graphics blur. Evaluating these practical experiments, visualization researchers could identify best practices and establish standards supporting other communities. With these steps, we hope that the visualization community advances scientific publications into visualization-enriched texts and other communities will adapt these practices. If successful, publications will get more effective as a medium for communicating scientific results and findings.

Acknowledgments

We thank Yasett Acurana, Tanja Blascheck, Corinna Vehlow, and Yuliya Volga for helping design and implementing the word-sized graphics listed in Section 4. Fabian Beck is indebted to the Baden-Württemberg Stiftung for the financial support of this research project within the Postdoctoral Fellowship for Leading Early Career Researchers.

References

[1] P. Ayres and J. Sweller, "The split-attention principle in multimedia learning," in *The Cambridge Handbook of Multimedia Learning*, R. E. Mayer, Ed. Cambridge University Press, 2005, pp. 135–146.

[2] J. Sweller, "Implications of cognitive load theory for multimedia learning," in *The Cambridge Handbook of Multimedia Learning*, R. E. Mayer, Ed. Cambridge University Press, 2005, pp. 19–30.

[3] J. Sweller, J. van Merrienboer, and F. Paas, "Cognitive architecture and instructional design," *Educational Psychology Review*, vol. 10, no. 3, pp. 251–296, 1998.

[4] P. Ginns, "Integrating information: A meta-analysis of the spatial contiguity and temporal contiguity effects," *Learning and Instruction*, vol. 16, no. 6, pp. 511–525, 2006.

[5] E. R. Tufte, *Beautiful Evidence*, 1st ed. Graphics Press, 2006.

[6] P. Goffin, J. Boy, W. Willett, and P. Isenberg, "An exploratory study of word-scale graphics in data-rich text documents," *IEEE Transactions on Visualization and Computer Graphics*, 2016. [Online]. Available: https://doi.org/10.1109/TVCG.2016.2618797

[7] L. Lebduska, "Emoji, emoji, what for art thou?" *Harlot: A Revealing Look at the Arts of Persuasion*, vol. 1, no. 12, 2014.

[8] R. Borgo, J. Kehrer, D. H. S. Chung, E. Maguire, R. S. Laramee, H. Hauser, M. Ward, and M. Chen, "Glyph-based visualization: Foundations, design guidelines, techniques and applications," in *Eurographics State of the Art Reports*. Eurographics Association, 2013, pp. 39–63.

[9] U. Brandes, B. Nick, B. Rockstroh, and A. Steffen, "Gestaltlines," *Computer Graphics Forum*, vol. 32, no. 3pt2, pp. 171–180, 2013.

[10] P. Goffin, W. Willett, A. Bezerianos, and P. Isenberg, "Exploring the effect of word-scale visualizations on reading behavior," in *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 2015, pp. 1827–1832.

[11] P. Goffin, W. Willett, J.-D. Fekete, and P. Isenberg, "Exploring the placement and design of word-scale visualizations," *IEEE Transactions on Visualization and Computer Graphics*, vol. 20, no. 12, pp. 2291–2300, 2014.

[12] F. Beck, S. Koch, and D. Weiskopf, "Visual analysis and dissemination of scientific literature collections with SurVis," *IEEE Transactions on Visualization and Computer Graphics*, vol. 22, no. 1, pp. 180–189, 2016.

[13] E. Adar, D. S. Weld, B. N. Bershad, and S. S. Gribble, "Why we search: visualizing and predicting user behavior," in *Proceedings of the 16th International Conference on World Wide Web*. ACM, 2007, pp. 161–170.

[14] A. T. Baker, "Theoretical and empirical studies of software development's role as a design discipline," Ph.D. dissertation, California State University at Long Beach, 2010.

[15] F. Beck, B. Dit, J. Velasco-Madden, D. Weiskopf, and D. Poshyvanyk, "Rethinking user interfaces for feature location," in *Proceedings of the 2015 IEEE 23rd International Conference on Program Comprehension*. IEEE, 2015, pp. 151–162.

[16] B. C. Boehmke, "Grabbing the air force by the tail: Applying strategic cost analytics to understand and manage indirect cost behavior," Ph.D. dissertation, Air Force Institute of Technology, 2015.

[17] J. Bryer and L. Daniels, "Measuring all students: An alternative method for retention and completion rates," in *Proceedings of the North East Association of Institutional Research Annual Conference*, 2011.

[18] A. Caglayan, M. Toothaker, D. Drapeau, D. Burke, and G. Eaton, "Behavioral analysis of botnets for threat intelligence," *Information Systems and e-Business Management*, vol. 10, no. 4, pp. 491–519, 2012.

[19] K. K. Chui, P. Webb, R. M. Russell, and E. N. Naumova, "Geographic variations and temporal trends of Salmonella-associated hospitalization in the US elderly, 1991-2004: A time series analysis of the impact of HACCP regulation," *BMC Public Health*, vol. 9, no. 1, pp. 1–10, 2009.

[20] J. J. Duda, M. M. Beirne, K. Larsen, D. Barry, K. Stenberg, and M. L. McHenry, "Aquatic ecology of the Elwha River estuary prior to dam removal," in *Coastal Habitats of the Elwha River Washington—Biological and Physical Patterns and Processes Prior to Dam Removal. USGS Scientific Investigations Report*, 2011, vol. 5120, pp. 175–223.

[21] D. Garrick, S. M. Whitten, and A. Coggan, "Understanding the evolution and performance of water markets and allocation policy: A transaction costs analysis framework," *Ecological Economics*, vol. 88, pp. 195–205, 2013.

[22] A. Gillespie, H. Sanei, A. Diochon, B. Ellert, T. Regier, D. Chevrier, J. Dynes, C. Tarnocai, and E. Gregorich, "Perennially and annually frozen soil carbon differ in their susceptibility to decomposition: analysis of subarctic earth hummocks by bioassay, XANES and pyrolysis," *Soil Biology and Biochemistry*, vol. 68, pp. 106–116, 2014.

[23] K. A. Janes, H. C. Reinhardt, and M. B. Yaffe, "Cytokine-induced signaling networks prioritize dynamic range over signal strength," *Cell*, vol. 135, no. 2, pp. 343–354, 2008.

[24] N. P. Johnson, D. T. Johnson, R. L. Kirkeeide, C. Berry, B. De Bruyne, W. F. Fearon, K. G. Oldroyd, N. H. Pijls, and K. L. Gould, "Repeatability of fractional flow reserve despite variations in systemic and coronary hemodynamics," *JACC: Cardiovascular Interventions*, vol. 8, no. 8, pp. 1018–1027, 2015.

[25] A. Jordan, "Evaluating and estimating the WCET criticality metric," in *Proceedings of the 11th Workshop on Optimizations for DSP and Embedded Systems*. ACM, 2014, pp. 11–18.

[26] S. Kavalanekar, B. Worthington, Q. Zhang, and V. Sharda, "Characterization of storage workload traces from production Windows servers," in *Proceedings of the IEEE International Symposium on Workload Characterization*. IEEE, 2008, pp. 119–128.

[27] I. Lang, E. Gardener, F. A. Huppert, and D. Melzer, "Was John Reid right? Smoking, class, and pleasure: a population-based cohort study in England," *Public Health*, vol. 121, no. 7, pp. 518–524, 2007.

[28] J. S. Lee and E. N. Waithaka, "The intersections of marginalized social identities in the transition to adulthood: A demographic profile," *Emerging Adulthood*, 2016. [Online]. Available: https://doi.org/10.1177/2167696816659021

[29] M. Lyra, D. Clarke, H. Morgan, J. Reffin, and D. Weir, "High value media monitoring with machine learning," *KI – Künstliche Intelligenz*, vol. 27, no. 3, pp. 255–265, 2013.
[30] J. R. Mahan, A. W. Young, and P. Payton, "Continuously monitored canopy temperature as a proxy for plant water status," *American Journal of Plant Sciences*, vol. 6, no. 14, p. 2287, 2015.
[31] I. Milligan, "Illusionary order: Online databases, optical character recognition, and Canadian history, 1997-2010," *Canadian Historical Review*, vol. 94, no. 4, pp. 540–569, 2013.
[32] S. Neuhaus and T. Zimmermann, "Security trend analysis with CVE topic models," in *Proceedings of the 21st International Symposium on Software Reliability Engineering*. IEEE, 2010, pp. 111–120.
[33] F. Peterson, J. Sexton, and K. Lajtha, "Scaling litter fall in complex terrain: A study from the western Cascades Range, Oregon," *Forest Ecology and Management*, vol. 306, pp. 118–127, 2013.
[34] C. Potts, "On the negativity of negation," in *Semantics and Linguistic Theory*, vol. 20, 2010, pp. 636–659.
[35] S. J. Schultheiss, M.-C. Münch, G. D. Andreeva, and G. Rätsch, "Persistence and availability of web services in computational biology," *PloS one*, vol. 6, no. 9, p. e24914, 2011.
[36] T. Torsvik, B. Lillebo, and G. Mikkelsen, "Presentation of clinical laboratory results: an experimental comparison of four visualization techniques," *Journal of the American Medical Informatics Association*, vol. 20, no. 2, pp. 325–331, 2013.
[37] K. Willadsen and J. Wiles, "Robustness and state-space structure of boolean gene regulatory models," *Journal of Theoretical Biology*, vol. 249, no. 4, pp. 749–765, 2007.
[38] C.-S. Wu, "Designing tangible tabletop interactions to support the fitting process in modeling biological systems," Ph.D. dissertation, Georgia Institute of Technology, 2012.
[39] A. T. Ying and M. P. Robillard, "Selection and presentation practices for code example summarization," in *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering*. ACM, 2014, pp. 460–471.
[40] U. Brandes and B. Nick, "Asymmetric relations in longitudinal social networks," *IEEE Transactions on Visualization and Computer Graphics*, vol. 17, no. 12, pp. 2283–2290, 2011.
[41] L. D. Frishberg, "Interactive sparklines: A dynamic display of quantitative information," in *CHI'11 Extended Abstracts on Human Factors in Computing Systems*. ACM, 2011, pp. 589–604.
[42] A. Kachkaev, J. Wood, and J. Dykes, "Glyphs for exploring crowd-sourced subjective survey classification," *Computer Graphics Forum*, vol. 33, no. 3, pp. 311–320, 2014.
[43] B. Lee, N. Henry Riche, A. K. Karlson, and S. Carpendale, "SparkClouds: Visualizing trends in tag clouds," *IEEE Transactions on Visualization and Computer Graphics*, vol. 16, no. 6, pp. 1182–1189, 2010.
[44] H. Song, B. Lee, B. Kim, and J. Seo, "DiffMatrix: Matrix-based interactive visualization for comparing temporal trends," in *EuroVis - Short Papers*. The Eurographics Association, 2012, pp. 103–107.
[45] C. Turkay, A. Slingsby, H. Hauser, J. Wood, and J. Dykes, "Attribute signatures: Dynamic visual summaries for analyzing multivariate geographical data," *IEEE Transactions on Visualization and Computer Graphics*, vol. 20, no. 12, pp. 2033–2042, 2014.
[46] W. Aigner, A. Rind, and S. Hoffmann, "Comparative evaluation of an interactive time-series visualization that combines quantitative data with qualitative abstractions," *Computer Graphics Forum*, vol. 31, no. 3, pp. 995–1004, 2012.
[47] D. T. Bauer, S. Guerlain, and P. J. Brown, "The design and evaluation of a graphical display for laboratory data," *Journal of the American Medical Informatics Association*, vol. 17, no. 4, pp. 416–424, 2010.
[48] L. M. Parsons and D. Tinkelman, "Testing the feasibility of small multiples of sparklines to display semimonthly income statement data," *International Journal of Accounting Information Systems*, vol. 14, no. 1, pp. 58–76, 2013.
[49] P. Goffin, W. Willett, J.-D. Fekete, and P. Isenberg, "Design considerations for enhancing word-scale visualizations with interaction," in *Proceedings of IEEE VIS 2015 Posters*, 2015.
[50] J. Heer and M. Agrawala, "Multi-scale banking to 45 degrees," *IEEE Transactions on Visualization and Computer Graphics*, vol. 12, no. 5, pp. 701–708, 2006.
[51] J. Parnow, "Micro Visualizations," Master's thesis, Potsdam University of Applied Sciences, 2015.
[52] J. Parnow and M. Dörk, "Micro Visualizations: Data-driven typography and graphical text enhancement," in *Proceedings of IEEE VIS 2015 Posters*, 2015.
[53] J. Oh, "Text visualization of song lyrics," Course cs448b report at Center for Computer Research in Music and Acoustics, Stanford University, 2010. [Online]. Available: https://ccrma.stanford.edu/~jieun5/cs448b/final/Oh_final.pdf
[54] D. Tinkelman, "Increasing the transparency and information content of financial statements using sparklines," Available at SSRN 1325998, 2009.
[55] F. Beck, T. Blascheck, T. Ertl, and D. Weiskopf, "Word-sized eye tracking visualizations," in *Eye Tracking and Visualization. Foundations, Techniques, and Applications (ETVIS 2015)*, M. Burch, L. Chuang, B. Fisher, A. Schmidt, and D. Weiskopf, Eds. Springer, 2016, pp. 113–128.
[56] J. H. Goldberg and J. I. Helfman, "Visual scanpath representation," in *Proceedings of the 2010 Symposium on Eye-Tracking Research & Applications*. ACM, 2010, pp. 203–210.
[57] C. Perin, R. Vuillemot, and J.-D. Fekete, "SoccerStories: A kick-off for visual soccer analysis," *IEEE Transactions on Visualization and Computer Graphics*, vol. 19, no. 12, pp. 2506–2515, 2013.
[58] M. Burch, T. Munz, F. Beck, and D. Weiskopf, "Visualizing work processes in software engineering with Developer Rivers," in *Proceedings of the 3rd Working Conference on Software Visualization*. IEEE, 2015, pp. 116–124.
[59] H. Barsnes, M. Vaudel, and L. Martens, "JSparklines: Making tabular proteomics data come alive," *Proteomics*, vol. 15, no. 8, pp. 1428–1431, 2015.
[60] A. Kowarik, B. Meindl, and M. Templ, "sparkTable: Generating graphical tables for websites and documents with R," *The R Journal*, vol. 7, no. 1, pp. 24–37, 2014.
[61] M. Templ, "Correlation between indicators over time in thematic maps," *Austrian Journal of Statistics*, vol. 41, no. 1, pp. 67–79, 2016.
[62] R. P. Radecki and M. A. Medow, "Cognitive debiasing through sparklines in clinical data displays," in *Proceedings of the 2007 AMIA Symposium*, vol. 11, 2007, p. 1085.
[63] F. Beck, Y. Acurana, T. Blascheck, R. Netzel, and D. Weiskopf, "An expert evaluation of word-sized visualizations for analyzing eye movement data," in *Proceedings of ETVIS 2016*. IEEE, 2016.
[64] C. Vehlow, F. Beck, and D. Weiskopf, "Visualizing group structures in graphs: a survey," *Computer Graphics Forum*, 2016. [Online]. Available: https://doi.org/10.1111/cgf.12872
[65] K. Kucher and A. Kerren, "Text visualization techniques: Taxonomy, visual survey, and community insights," in *Proceedings of the 2015 IEEE Pacific Visualization Symposium*. IEEE, 2015, pp. 117–121.
[66] T. Blascheck, K. Kurzhals, M. Raschke, M. Burch, D. Weiskopf, and T. Ertl, "State-of-the-art of visualization for eye tracking data," in *EuroVis - STARs*. Eurographics Association, 2014, pp. 63–82.
[67] F. Beck, M. Burch, S. Diehl, and D. Weiskopf, "A taxonomy and survey of dynamic graph visualization," *Computer Graphics Forum*, 2016. [Online]. Available: htpps://doi.org/10.1111/cgf.12791

Fabian Beck is assistant professor at the University of Duisburg-Essen, Germany. He received the Dr. rer. nat. (PhD) degree in computer science from the University of Trier, Germany in 2013. His research focuses on methods for visualizing and comparing large and dynamic graphs and hierarchies, often in the context of software systems and their evolution. He also investigates visual analytics systems and the integration of visualizations into text.

Daniel Weiskopf is a professor at VISUS, University of Stuttgart, Germany. He received his Dr. rer. nat. (PhD) degree in physics from the University of Tübingen, Germany (2001), and the Habilitation degree in computer science at the University of Stuttgart, Germany (2005). His research interests include all areas of visualization, visual analytics, GPU methods, perception-oriented computer graphics, and special and general relativity.