Burger Charts: A Quantitative Display of Set Intersections

P. Riehmann[2], D. Kiesel[1], J. König[1], and B. Froehlich[1]

[1]Virtual Reality and Visualization, Bauhaus-Universität Weimar, Germany
[2] Department of Computer Science and Informatics, Jönköping University, Sweden

Abstract

We present Burger Charts, a quantitative display of set intersections that stack the contributing sets on top of each other. Sets are represented as horizontal layers. The size of each set intersection is encoded by the width of its layer section, forming a vertically stacked, burger-like visual representation. A visual skewer maintains the unity of the burger by bridging gaps of those set layers that do not contribute to the intersection. The color coding of the sets emphasizes which set contributes to which intersection. We use Burger Charts to visualize and analyze keyword co-occurrences of an argument search engine. They support a quantitative discourse analysis by providing insights into the distribution of sets of keyword occurrences. Users can interactively explore keywords of online discourses on controversial topics, identify prevalent keyword co-occurrences, and even uncover overlooked perspectives.

CCS Concepts

• ***Information systems*** → *Search interfaces;* • ***Human-centered computing*** → ***Visual analytics;*** *Visualization systems and tools;*

1. Introduction

Sets and set intersections occur in many domains. Sets can be formed in various ways. In our research, we assign texts to one or multiple sets based on whether a certain keyword appears in the text. In such a case, the domain-specific research questions often translate to set analysis questions such as: (1) How many intersections of sets (meaning a combination of keywords) exist? (2) What intersections does a set contribute to? (3) How large is an intersection (meaning how often does a certain combination exist), particularly in comparison to all other combinations? To provide answers to these questions, set visualization techniques such as Parallel Sets [KBH06], Mosaic Plots [HK81, Fri94], UpSet [LGS*14], and Linear Diagrams [RSC15] can be used. These techniques use different marks and channels for encoding sets and set intersections. Parallel Sets rely on the connections of ribbons and their width. Mosaic Plots use the position and size of rectangles. UpSet highlights circles in rows, and Linear Diagrams employ aligned line segments. By design, all have difficulties answering question 3 about (relative) quantities of different intersections directly – either at all (Linear Diagrams), for more than two or three attributes (Mosaic Plots), without the help of external or attached visualizations (UpSet), or without interactive highlighting (Parallel Sets).

Thus, we designed the Burger Chart to provide a better quantitative overview of the distribution of intersection sizes at first glance (see Figure 1). The basic idea is that each existing set is shown as a horizontal layer. The layers of sets that contribute to a particular intersection are stacked in a burger-like manner in the vertical direction. The Gestalt law of proximity visually ties the burger's layers to each other. A visual skewer bridges gaps of set layers that do not contribute to this intersection and thus maintains connectedness and unity of the burger. In order to separate a single burger from others, the Gestalt law of closure is applied by using rounded corners at the top and bottom layer – like a real burger bun. That way, the space between the burgers can be minimal. An additional

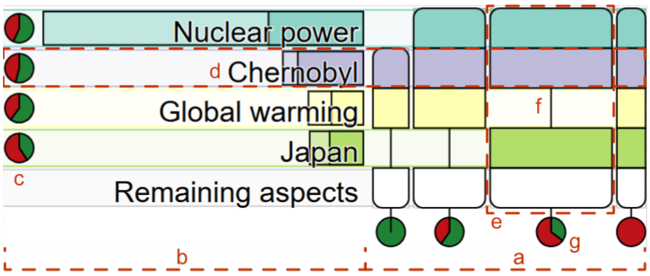

Figure 1: *Four sample Burger Charts (a). The left part shows the aspect keywords that determine the sets (b) with pro/con distributions (c). The four Burgers (a) intersect different aspect keyword combinations, yet all include "Chernobyl" (d). The marked Burger (e) contains all but "Global warming," which is bridged by a skewer (f). The width shows that this intersection contains the most contributions of the four. Attached to each Burger at its skewer is a pie chart expressing the respective pro/con distribution (g).*

DOI: 10.2312/evs.20251084

color coding of the sets emphasizes which sets contribute to which intersection. But even more importantly, the width of a burger encodes the intersection's size to provide a clear comparison of the relative quantities of all visible intersections.

We demonstrate our technique by analyzing controversial topics with the argument search engine *args.me* that retrieves discourse contributions for given topic queries. Each contribution is characterized by keywords that are called aspects [WPAK*17, AWK*18] – e.g. "fossil fuels," "pollution,", and "Chernobyl" are keywords for contributions of the query "nuclear energy" (see Figure 2). The discourse contributions are gathered from online debate forums, foremost debate.org, and usually contain a chain of arguments about different aspects of a certain topic. The original *args.me* interface just lists the high number of discourse contributions (for instance, 5 107 for "Death Penalty"), divided into pro and con stance, but does not further consider the aspects, not even for grouping them. Yet, the aspects could support the user in exploring the topics by providing an overview of frequent subtopics or related topics in the discourse contributions and answering research questions such as: "Which combination (intersection) of aspects is most heatedly discussed,", "What are the most discussed aspects in renewable energies," or "Does the discussion focus on many or only a few aspects?"

Concerning these questions, we developed the *Quantitative Discourse Analyzer* featuring the Burger Chart technique for visualizing, analyzing, and managing co-occurrences of aspects as a quantitative display of set intersections. An aspect set includes the discourse contributions (texts) that contain the aspect's keyword. An intersection of aspect sets contains the discourse contributions where the respective keywords co-occur. Our design allows the user to get an overview of the most frequent aspects, their interrelations, and their stance regarding the search query. Furthermore, the view can be customized by adding, removing, and grouping aspects according to the information need. So, the Burger Chart allows for a quick assessment of which aspects are contained in an intersection and encodes the size of the intersection directly. This conveys an immediate impression of the proportions between the intersections, naturally emphasizing frequent ones. Moreover, the fusion of intersections allows for a detailed analysis of the frequencies of keyword co-occurrences. Overall, the system enables quantitative analysis of the frequencies and co-occurrences of aspects compiled from moderated and curated real-world discourse contributions. This reveals insights about the individual importance of the aspects and their relations that neither a list interface of a search engine nor even a text written by an LLM can provide.

2. Related Work

Other set-based visualizations (see Alsallakh's survey [AMA*16]) have been presented in the past. Riche [RD10] proposed enhanced Euler diagrams for expressing intersections by splitting sets into connected subsets or duplicating items. Concerning the number of set intersections and the number of discourse contributions to expect, they are not applicable. Radial Sets [AAMH13] depict sets as segments of a ring-like shape where the central space is used to visualize the intersecting sets as count-dependent circles. Yet, the circles' positions are not clearly aligned making it necessary to rely on interaction to identify the corresponding sets. A clear alignment is given by Linear Diagrams [RSC15] where each set is represented as one or more line segments, with all sets drawn in parallel. Where lines overlap, the corresponding intersection of sets contains an element that is not in any of the remaining sets. Moreover, between them, the overlaps represent all of the non-empty set intersections. While this technique is suitable for tasks concerning sets, the intersections are not as clearly separated. Contrary to that, UpSet [LGS*14] focuses on the set intersections showing them as dots and dot connections in a row augmented with associated visualizations and data. Our approach combines the advantages of both. Rainbow Boxes [LBCF17] assign unique colors to set intersections by calculating the mean (as RGB) of the distinct colors of the parent sets but do not encode set or intersection sizes. Onset [SMDS14] introduces specialized layouts and aggregation strategies for visualizing up to thousands of sets. The separation of sets into different graphical elements does not provide an overview of the interrelations between sets without interaction.

Dimensionality reduction techniques such as MDS, PCA, and t-SNE (see surveys and comparisons [VDMPVDH*09, CG15, AHT20]) often based on curated topics or a prior topic modeling step (like LDA [BNJ03]) are used to express relations between documents and topics (aspects). This inspired many visualizations and embeddings in 2D such as galaxies [WTP*95], as visual tool for method analysis [NA19], as navigator-based approach [CB21] or as text flows or topic evolvement [CLT*11, GHN12, TQW*24]. However, the output of these reduction techniques as 2D embedding relies on the similarity between documents and visual clusters in 2D. So, their affiliation to an individual aspect is not clear. Generalized barycentric coordinates [MBLD02] improve the association between documents and aspects, since the location of a document in the aspect space embedding is defined as a weighted linear combination of a regular polygon whose vertices represent the top-ranking aspects. Positions in a 2D polygonal topic space do not have unique generalized barycentric coordinates. So, very different topic combinations can result in similar or even the same positions resulting in misinterpretations. Additionally, overplotting and clutter occur the larger the number of contributions. There are approaches to mitigate this by transforming the aspect space (such as Cheng [CM15] and Zanabria [ZNGN16]) or by introducing glyphs instead of points, whose shapes resemble the topic space and thus use spikes to encode the affinity towards the individual topics. [RKKF18, KRF*18, AWK*18]. This reduces ambiguity towards the topics but doesn't help with overplotting and clutter. To express the affiliation of documents/texts towards topics/aspects matrix-like visualization such as Termite [CMH12], and Serendip [AKV*14] can be also very useful.

3. Visual Design and Interaction

The Quantitative Discourse Analyzer enables the user to gain quantitative insight into aspect occurrences in discussion contributions, to understand co-occurrences of aspects, and, eventually, to view the aspects' keywords in the context of the original texts. It consists of two parts: The intersection view provides information about the aspects and their co-occurrences; the aspect overview shows all

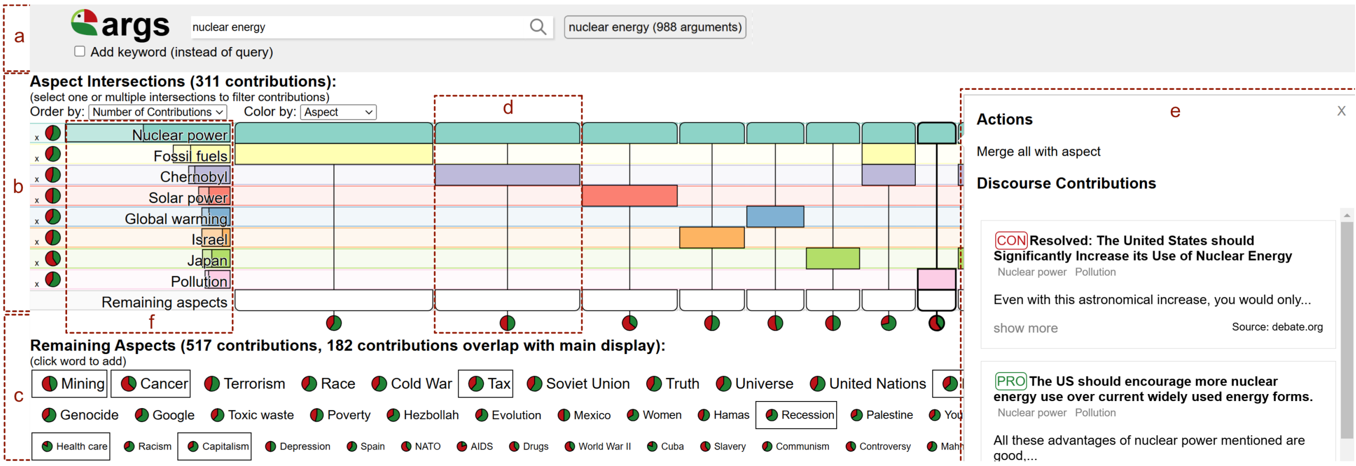

Figure 2: *The system consists of the search field (a) and two aspect visualizations (b and c). The intersection view (b) shows the aspects as sets in color-coded rows and the set intersections as columns of connected aspect combinations (d). The width of an intersection encodes the number of discourse contributions that contain this aspect combination. The intersection on the right has been selected to read the actual discourse contributions shown next to it (e). A bar behind each aspect's keyword shows its frequency relative to the most frequent keyword; The darker bar encodes the fraction of discourse contributions with this keyword shown in the intersection view (f). The remaining aspect overview (c) shows all the aspects' keywords that would not fit into the intersection view.*

the remaining aspects of the result set. A simple click adds aspects from the aspect overview to the intersection view.

Design: The intersection view initially shows the eight most frequent aspects in descending order, one aspect per row (see Figure 1). The burger columns of the intersection view represent the intersections between the aspects. The width of a column encodes the number of discourse contributions in the represented intersection. The distinct shape, a rectangle with rounded corners that spans vertically across the rows, resembles a burger. Each involved aspect adds a layer in the aspect's color to the intersection representative. Aspect's rows not part of the intersection will be crossed skewer-like to provide visual continuation and connectedness for the intersection. The top and bottom layers are rounded to provide an enclosed appearance for each intersection by bridging the missing layers via the Gestalt principle of closure. Clicking on an intersection opens a detail window containing additional actions, such as merging or splitting intersections, and the list of all discourse contributions in the intersection.

Pro/Con: We use pie charts to express relationships between the number of pro versus con discourse contributions for each aspect (attached to the left) and for each intersection (connected to it below its bottom layer). The pie chart shows how balanced (or not) the discussion within the sets and intersections is, where pro or con dominates, or where to find discourse contributions backing one's own opinion. Kosara [Kos19a, Kos19b] showed the advantages of pie charts for part-to-whole comparisons, which is exactly what we do by showing the ratio between the pro and con conditions.

General Layout: By default, we start with eight aspects, ensuring a good balance between usefulness and readability. The more aspects are included in the intersection view, the more fragmented the display can become since smaller intersections occur. The aspects that appear in the discourse contributions yet not in the intersection view are summarized in the bottom white row. On hover, the contained aspects are highlighted in the aspect overview below the intersection view. The intersection display is ordered by size from left to right. This way, the potentially most important aspect combinations can quickly be compared and the result looks tidy.

Orientation: Research on Linear Diagrams [RSC15] showed that orientation does not matter in terms of speed or accuracy. Thus, the orientation of the intersection view is entirely up to layout considerations. A vertical orientation – having the aspect sets as columns and the intersections as rows – grants more space to the intersections thus allowing for either more intersections to be viewed at once or larger intersection rectangles. However, the height of the intersection views would then push the aspect overview beyond screen space, thus cutting the connection between the two parts of the visualization, which is why we opted for a horizontal layout.

Color: The color model is either aspect-set-centric or intersection-centric. Emphasizing aspect membership is key to assessing how many intersections a specific aspect is part of. We use distinct colors for each aspect in the respective intersection layers and (toned down) for the row background. The visual cues for intersections rely on containment and connection lines. An intersection-centric alternative is a coloring as suggested for Rainbow Boxes [LBCF17] showing all rows of an intersection in the same color, calculated as average color (in RGB space) of the individual aspect's colors of the intersection (see Figure 3).

Interactions: Tailored interactions support a deeper analysis of the search results. Hovering over an aspect's keyword highlights all intersection layers connected to the aspect. To manage the given aspect sets and shape the display to the user's needs, the **intersections** and **aspect sets** can be **reordered** via drag and drop. Addi-

Proceedings published by Eurographics - The European Association for Computer Graphics.

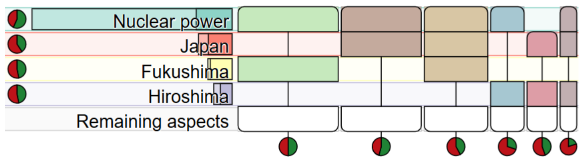

Figure 3: *The color scheme can be changed to have one color per intersection so that intersection-based tasks become easier.*

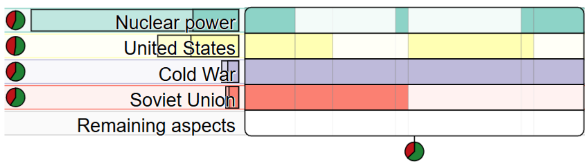

Figure 4: *An example of a fused intersection: All intersections containing the aspect "Cold War" have been fused and sorted by "United States" (yellow) and "Soviet Union" (red). It shows that in the context of "Nuclear energy" "Soviet Union" is mentioned in about 50% of the texts containing "Cold War" and in about 50% of those texts also "United States" occurs.*

tionally, **aspect sets** can also be **merged** via drag and drop such that the resulting row corresponds to both aspects' keywords, mimicking a logic or operation. This is helpful since different keywords that describe similar concepts or occur in the same context can be summarized into one aspect. **Keywords** can also be **added** to the intersection view out of the remaining aspect overview at the bottom or **removed** back to it. New **keywords** can be **created** through the add-keyword functionality of the text-input field.

Fused Intersections: Intersections can be merged into **fused intersections** via drag and drop or the "merge all with aspect" option of the detail window. The individual intersection layers (empty or not) are vertically stacked to form a single layer for each aspect. The compact display and sorting options of the fused intersection allow the relative quantification of keyword co-occurrences. Figure 4 shows a fused intersection based on all intersections containing the keyword "Cold War." According to the display, about 50% of texts containing "Cold War" also contain "Soviet Union" and 50% of those also contain "United States."

4. Use cases and User Feedback

Our use case throughout the paper is based on the search query "nuclear energy" which yields initial aspects that suggest three argumentative perspectives (see Figure 2): nuclear power in contrast to other energy sources like "fossil fuels" or "solar power" incidents with nuclear power in "Chernobyl" and "Japan," as well as an environmental perspective with "global warming" and "pollution." The keyword "Israel" does not seem to fit, but when reading the corresponding texts "Israel" is mentioned as owning a nuclear weapon or being threatened by them. In "Japan" actually, two incidents occurred that involved nuclear energy: The nuclear bomb in Hiroshima during World War II and the nuclear accident in Fukushima. Adding these as keywords lets the analyst determine the interrelation of those aspects (see Figure 3): "Fukushima" is mentioned more often than "Hiroshima." All intersections with "Hiroshima" tend to argue against nuclear energy, while the intersection of only "Fukushima" and "Nuclear Power" (the first one) contains a very balanced discussion. The aspect overview figures "Cold War" and "Soviet Union" as frequent keywords hinting at another perspective on nuclear energy. To investigate, we add both aspects to the intersection view and create a new aspect "United States" to represent both superpowers during that time. To improve the estimability of the co-occurrences, we create a fused intersection of all contributions containing "Cold War" (see Figure 4). Ordering according to "United States" and "Soviet Union" shows that in only half of the contributions with "Cold War" "Soviet Union" is mentioned. And only about 50% of those also figure "United States." Even though an event worth mentioning, it seems that the Cold War is not discussed in depth (or with other keywords) in the discourse contributions.

We gathered feedback from an expert of the development team of the *args.me* search engine who was not involved in any capacity during the design of the Burger Chart nor associated to development process of the Quantitative Discourse Analyzer. At first, the expert was introduced to the visualization concepts and interaction possibilities. Then, the expert freely explored the system and answered questions of interest. Adding custom keywords helped refine the display in the search for arguments regarding the query "education should be free," and helped to disambiguate from other related topics like whether education should be mandatory. When exploring the topic "abortion," the expert was surprised to find the keyword "cancer." A look at some sample texts revealed that abortion (allegedly, only according to people in the debate forum) is supposed to increase the chance of getting breast cancer, a consequence the expert was not aware of yet. One reported drawback was that the tool could not perform logic and operations between aspects since filtering for multiple aspects at once is not implemented.

5. Conclusion and Future Work

We presented the Burger Charts as the key concept for the Quantitative Discourse Analyzer, a tool for analyzing aspect sets and set intersections in online discourse. It is an interactive display of the distribution of keywords as aspect sets and their co-occurrences as set intersections. The user quickly gains insights regarding the most frequent aspects discussed for the search query, finds typical aspect combinations via the set intersections, and can analyze the relative frequency of keyword co-occurrences through fused intersections. The possibility to add and remove aspects, order and group aspects and intersections, and the easy access to the textual content of the discourse contributions allow for a flexible and versatile workflow that supports the analyst in aspect-set- and intersection-based tasks.

Even though designed for the analysis of online discourse, the Burger Chart concept can be easily adapted to all kinds of data items that can be assigned to multiple sets.

Acknowledgments

This work was supported in part by the German Federal Ministry of Education and Research (BMBF) under grant 16DHBKI083 (project SKILL).

References

[AAMH13] Alsallakh B., Aigner W., Miksch S., Hauser H.: Radial sets: Interactive visual analysis of large overlapping sets. *IEEE Transactions on Visualization and Computer Graphics 19*, 12 (2013), 2496–2505. doi:10.1109/TVCG.2013.184. 2

[AHT20] Ayesha S., Hanif M. K., Talib R.: Overview and comparative study of dimensionality reduction techniques for high dimensional data. *Information Fusion 59* (2020), 44–58. doi:10.1016/j.inffus.2020.01.005. 2

[AKV*14] Alexander E., Kohlmann J., Valenza R., Witmore M., Gleicher M.: Serendip: Topic model-driven visual exploration of text corpora. In *2014 IEEE Conference on Visual Analytics Science and Technology (VAST)* (2014), pp. 173–182. doi:10.1109/VAST.2014.7042493. 2

[AMA*16] Alsallakh B., Micallef L., Aigner W., Hauser H., Miksch S., Rodgers P.: The State-of-the-Art of Set Visualization. *Comput. Graph. Forum 35*, 1 (Feb. 2016), 234–260. doi:10.1111/cgf.12722. 2

[AWK*18] Ajjour Y., Wachsmuth H., Kiesel D., Riehmann P., Fan F., Castiglia G., Adejoh R., Stein B., Froehlich B.: Visualization of the topic space of argument search results in args.me. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (2018). doi:10.18653/v1/D18-2011. 2

[BNJ03] Blei D. M., Ng A. Y., Jordan M. I.: Latent dirichlet allocation. *Journal of Machine Learning Research. 3* (Mar. 2003), 993–1022. 2

[CB21] Chaney A., Blei D.: Visualizing topic models. *Proceedings of the International AAAI Conference on Web and Social Media 6*, 1 (Aug. 2021), 419–422. doi:10.1609/icwsm.v6i1.14321. 2

[CG15] Cunningham J. P., Ghahramani Z.: Linear dimensionality reduction: survey, insights, and generalizations. *J. Mach. Learn. Res. 16*, 1 (Jan. 2015), 2859–2900. 2

[CLT*11] Cui W., Liu S., Tan L., Shi C., Song Y., Gao Z., Qu H., Tong X.: Textflow: Towards better understanding of evolving topics in text. *IEEE Transactions on Visualization and Computer Graphics 17*, 12 (2011), 2412–2421. doi:10.1109/TVCG.2011.239. 2

[CM15] Cheng S., Mueller K.: Improving the fidelity of contextual data layouts using a generalized barycentric coordinates framework. In *Visualization Symposium (PacificVis), 2015 IEEE Pacific* (2015), IEEE, pp. 295–302. doi:10.1109/PACIFICVIS.2015.7156390. 2

[CMH12] Chuang J., Manning C. D., Heer J.: Termite: visualization techniques for assessing textual topic models. In *Proceedings of the International Working Conference on Advanced Visual Interfaces* (New York, NY, USA, 2012), AVI '12, Association for Computing Machinery, p. 74–77. doi:10.1145/2254556.2254572. 2

[Fri94] Friendly M.: Mosaic displays for multi-way contingency tables. *Journal of the American Statistical Association 89*, 425 (1994), 190–200. doi:10.1080/01621459.1994.10476460. 1

[GHN12] Gansner E. R., Hu Y., North S.: Visualizing streaming text data with dynamic graphs and maps. In *International Symposium on Graph Drawing* (2012), Springer, pp. 439–450. URL: 10.1007/978-3-642-36763-2_39. 2

[HK81] Hartigan J. A., Kleiner B.: Mosaics for contingency tables. In *Computer Science and Statistics: Proceedings of the 13th Symposium on the Interface* (New York, NY, 1981), Eddy W. F., (Ed.), Springer US, pp. 268–273. 1

[KBH06] Kosara R., Bendix F., Hauser H.: Parallel sets: Interactive exploration and visual analysis of categorical data. *Transactions on Visualization and Computer Graphics (TVCG) 12*, 4 (2006), 558–568. doi:10.1109/TVCG.2006.76. 1

[Kos19a] Kosara R.: Evidence for area as the primary visual cue in pie charts. In *IEEE VIS Short Paper Proceedings* (2019). doi:10.2312/evs20191162. 3

[Kos19b] Kosara R.: The impact of distribution and chart type on part-to-whole comparisons. In *Short Paper Proceedings of the Eurographics/IEEE VGTC Symposium on Visualization (EuroVis)* (2019). doi:10.2312/evs20191162. 3

[KRF*18] Kiesel D., Riehmann P., Fan F., Ajjour Y., Wachsmuth H., Stein B., Froehlich B.: Improving barycentric embeddings of topics spaces. In *Proceedings of the 24th IEEE Information Visualization Conference* (Berlin, October 2018). 2

[LBCF17] Lamy J.-B., Berthelot H., Capron C., Favre M.: Rainbow boxes: A new technique for overlapping set visualization and two applications in the biomedical domain. *Journal of Visual Languages & Computing 43* (2017), 71–82. doi:10.1016/j.jvlc.2017.09.003. 2, 3

[LGS*14] Lex A., Gehlenborg N., Strobelt H., Vuillemot R., Pfister H.: Upset: Visualization of intersecting sets. *IEEE Transactions on Visualization and Computer Graphics 20*, 12 (2014), 1983–1992. doi:10.1109/TVCG.2014.2346248. 1, 2

[MBLD02] Meyer M., Barr A., Lee H., Desbrun M.: Generalized barycentric coordinates on irregular polygons. *Journal of graphics tools 7*, 1 (2002), 13–22. doi:10.1080/10867651.2002.10487551. 2

[NA19] Nonato L. G., Aupetit M.: Multidimensional projection for visual analytics: Linking techniques with distortions, tasks, and layout enrichment. *IEEE Transactions on Visualization and Computer Graphics 25*, 8 (2019), 2650–2673. doi:10.1109/TVCG.2018.2846735. 2

[RD10] Riche N. H., Dwyer T.: Untangling euler diagrams. *IEEE Transactions on Visualization and Computer Graphics 16*, 6 (2010), 1090–1099. doi:10.1109/TVCG.2010.210. 2

[RKKF18] Riehmann P., Kiesel D., Kohlhaas M., Froehlich B.: Visualizing a thinker's life. *IEEE Transactions on Visualization and Computer Graphics* (2018). doi:10.1109/TVCG.2018.2824822. 2

[RSC15] Rodgers P., Stapleton G., Chapman P.: Visualizing Sets with Linear Diagrams. *ACM Trans. Comput.-Hum. Interact. 22*, 6 (sep 2015). doi:10.1145/2810012. 1, 2, 3

[SMDS14] Sadana R., Major T., Dove A., Stasko J.: Onset: A visualization technique for large-scale binary set data. *IEEE Transactions on Visualization and Computer Graphics 20*, 12 (2014), 1993–2002. doi:10.1109/TVCG.2014.2346249. 2

[TQW*24] Tu Y., Qiu R., Wang Y.-S., Yen P.-Y., Shen H.-W.: Phrasemap: Attention-based keyphrases recommendation for information seeking. *IEEE Transactions on Visualization and Computer Graphics 30*, 3 (2024), 1787–1802. doi:10.1109/TVCG.2022.3225114. 2

[VDMPVDH*09] Van Der Maaten L., Postma E. O., Van Den Herik H. J., et al.: Dimensionality reduction: A comparative review. *Journal of machine learning research 10*, 66-71 (2009), 13. 2

[WPAK*17] Wachsmuth H., Potthast M., Al-Khatib K., Ajjour Y., Puschmann J., Qu J., Dorsch J., Morari V., Bevendorff J., Stein B.: Building an Argument Search Engine for the Web. In *Proceedings of the Fourth Workshop on Argument Mining (ArgMining 17)* (2017). doi:10.18653/v1/W17-5106. 2

[WTP*95] Wise J., Thomas J., Pennock K., Lantrip D., Pottier M., Schur A., Crow V.: Visualizing the non-visual: spatial analysis and interaction with information from text documents. In *Proceedings of Visualization 1995 Conference* (1995), pp. 51–58. doi:10.1109/INFVIS.1995.528686. 2

[ZNGN16] Zanabria G. G., Nonato L. G., Gomez-Nieto E.: istar (i*): An interactive star coordinates approach for high-dimensional data exploration. *Computers & Graphics 60* (2016), 107–118. doi:10.1016/j.cag.2016.08.007. 2