

From Clusters to Graphs – Toward a Scalable Viewing of News Videos

Nicolas Ruth, Bernhard Liebl and Manuel Burghardt

Computational Humanities Group, Institute for Computer Science, Leipzig University, Germany

Abstract

In this paper, we present a novel approach that combines density-based clustering and graph modeling to create a scalable viewing application for the exploration of similarity patterns in news videos. Unlike most existing video analysis tools that focus on individual videos, our approach allows for an overview of a larger collection of videos, which can be further examined based on their connections or communities. By utilizing scalable reading, specific subgraphs can be selected from the overview and their respective clusters can be explored in more detail on the video frame level.

Keywords

scalable viewing, hdbscan clustering, community detection, graph visualization

1. Introduction: Video Analytics and Scalable Viewing

In digital humanities, *distant reading* has become a widely-used method, offering a macroanalytical alternative to the *close reading* of texts. This methodological shift has spurred considerable interest among scholars seeking novel ways to analyze textual material, ultimately leading to the proposal of a hybrid solution called *scalable reading* [23]. Visualization has always been a cornerstone of the distant reading methodology, exemplified by Franco Moretti’s influential work, which proposes graphs, maps and trees as fundamental text visualization techniques [14]. A number of scholars have since explored different ways to visualize quantitative patterns in large text corpora [7, 6]. The success of text-focused distant reading has further encouraged the extension of this paradigm into the realm of visual media, aptly termed *distant viewing* [2]. Visualization also plays a critical role in this domain, offering the means to unveil intricate patterns and interdependencies within large collections of visual media. While many of the fundamental concepts presented here draw inspiration from information visualization [22], interface design [20] and visual analytics [8], there are also some more specific approaches. Notably, the idea of “cultural analytics” [12] has been introduced to encompass the visualization of visual material, while “visual movie analytics” [9] has been suggested specifically for videos and movies. A more recent development in this domain has been the blend of close and distant


CHR 2023: Computational Humanities Research Conference, December 6 – 8, 2023, Paris, France

✉ nicolas.ruth@studserv.uni-leipzig.de (N. Ruth); bernhard.liebl@uni-leipzig.de (B. Liebl); burghardt@informatik.uni-leipzig.de (M. Burghardt)

🌐 <https://ch.uni-leipzig.de/> (N. Ruth); <https://ch.uni-leipzig.de/> (B. Liebl); <https://ch.uni-leipzig.de/> (M. Burghardt)

🆔 0000-0002-1126-6280 (N. Ruth); 0000-0002-8593-400X (B. Liebl); 0000-0003-1354-9089 (M. Burghardt)

© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

analysis perspectives for video material, termed “scalable viewing” [3, 16]. However, for video material, a significant challenge lies in providing a comprehensive overview from a distant perspective in a dynamic medium.

In this paper, we propose a novel approach to address this challenge by visualizing relationships and patterns of similarity within a video collection. The basis for our research stems from the ongoing *FakeNarratives*¹ project, in which we investigate the use of narrative strategies for the purpose of disinformation in German news videos [21]. We build on the *Zoetrope* prototype [11, 10], a tool developed in the project’s early phase, initially designed for the analysis of individual news videos. *Zoetrope* relies on a complex multimodal information extraction pipeline from which we have acquired CLIP embeddings [18] for a sample of 117 German news videos from “Tagesschau”, including all the videos from 01.01.2022 to 14.03.2022.

In the following, we present an approach that allows researchers to interactively navigate extensive video collections and to explore underlying patterns that are latent in the CLIP embeddings. Our approach opens up novel opportunities for scalable viewing of visual media in computational humanities research. The suggested approach enhances the existing landscape of video analysis tools such as the *Distant Viewing Toolkit* [1] or *VIAN* [5], which are mostly focused on the analysis of single videos.² In a sense our interactive visualization can be compared to recent cultural analytics tools, such as *PixPlot*³ or the *Collection Space Navigator* [15], which visualize patterns in large collections of images. However, we extend these existing tools as we propose a two-fold approach that combines initial clustering of similar images and a 3D network visualization of similar image clusters. Also, our approach is optimized for being used with video frames rather than other types of imagery. Please note that we describe a novel analytical workflow rather than a ready-to-use tool.⁴ While the use case we present is based on the analysis of news videos, the workflow can also be adapted for other collections of video material.

2. From Clusters to Graphs – Conceptual Approach

The employed methodology aims to summarize the visual content of news formats and to construct an exploratory visualization to investigate their connections and interrelationships. In the following sections, we describe our conceptual approach in some more detail and also provide information about the utilized clustering and graph algorithms and their parameterization.

¹Project website and further information: <https://fakenarratives.github.io/>

²For a comprehensive overview of similar video analysis tools see the survey paper by Pustu-Iren, Sittel, Mauer, Bulgakowa, and Ewerth [17].

³Yale DH Lab – PixPlot: <https://github.com/YaleDHLab/pix-plot>

⁴The full code is available via <https://github.com/Nicolas-le/from-clusters-to-graphs>. Due to copyright issues, we currently cannot share a demo visualization of the news video use case. However, to get a basic idea of what the interactive visualization actually looks like, you will also find a short demo video (2:05 min) in the above repository.

2.1. Feature Extraction via CLIP

First, we are concerned with the extraction of semantic elements from the video material [2] by means of CLIP embeddings, which were calculated for every frame every 0.5 seconds for each of the 117 news videos in the dataset, resulting in a total of 300,029 embedding vectors. Contrastive Language-Image Pre-Training (CLIP) is a neural network that was trained on 400 million image-text pairs with the primary objective to combine the textual description of an image and the image itself within a common vector space. As a result, CLIP embeddings have the ability to represent the content description of an image in a high-dimensional vector space. Therefore, embedding vectors with high similarity indicate a resemblance in the semantic content of the corresponding images [18].

2.2. Dimensionality Reduction and Density-Based Clustering

Next we aggregate similar content across videos via a clustering of the embedding vectors. To reduce the dimensionality of the embeddings, the data was standardized using the scikit-learn standard scaler and then followed by a Principal Component Analysis (PCA), to retain only 30 principal components.⁵ The parameterization of the PCA was the result of a mutual adjustment with the parameters of the clustering algorithm with the focus on the interpretability of the visualization and had mainly the goal to reduce the computation time of the clustering. In subsequent extensions of the procedure, this part can be extended by systematic determination of the principal components by, for example, a threshold of the explained variance and by tests with further dimension reduction procedures. In the early stages of the project, we experimented with KMeans clustering. However, this approach had certain limitations, including the clustering of noise and the requirement to define a fixed number of clusters. KMeans also showed some problems in dealing with outliers. Following that, our transition to a density-based clustering approach led us to using the HDBScan algorithm. The *Hierarchical Density-Based Spatial Clustering of Applications with Noise* (HDBScan) extends DBScan with a complete clustering hierarchy composed of all possible density-based clusters. It thus inherits the advantages of density-based clustering, among others the recognition of noise and the non-static number of clusters, and improves them by the ability to detect clusters of varying densities [4]. The algorithm can be parameterized in numerous manners. For instance, the Euclidean distance has been designated as the distance metric. However, a critical parameter to consider pertains to the minimum number of data points `min_cluster_size` required to constitute a cluster. There is typically no upper limit, so the algorithm is not likely to split bigger clusters into smaller sub-clusters, which can lead to clusters with significantly varying sizes. The optimization of `min_cluster_size` poses challenges due to a trade-off: On one hand, lower parameter values enable the algorithm to detect fine-grained clean clusters, often caused by consecutive video frames. On the other hand, higher values may lead to the complete loss of smaller clusters. Another trade-off is introduced by the second parameter, the *minimum number of neighbors* for a point to be a core point. Lowering the parameter values can result in the formation of larger clusters, which, in turn, amplifies the issue of

⁵Scikit-learn standard scaler: <https://scikit-learn.org/stable/modules/generated/sklearn.preprocessing.StandardScaler.html>

semantically *impure* clusters, which refers to clusters in which the topics exhibited within the frames of the cluster manifest considerable heterogeneity. Semantically *pure* clusters are favored by higher values, but they increase the detection of noise and encourage the focus on overly dense clusters created by successive frames or clusters with frequently occurring items, like recurring images. After conducting extensive experiments to compare the results and to assess usability in the visualization process, the `min_cluster_size` was set to 100, to detect fine-grained clusters. The minimum threshold for a core point was set to default, that is the same as the `min_cluster_size`. However, the fine-tuning of the parameters needs some more consideration in future work. In our current approach, the detection of small clusters was given greater weight, which resulted in the detection of altogether 231 clusters of varying size. The deliberate emphasis on smaller, semantically more distinctive clusters also favors the incidental classification of substantial quantities of data points as noise. To counteract the focus on too similar data points, the points previously clustered as noise were added again to the clustering, using a function in the HDBScan implementation of [13] to assign new data points into the computed clustering, which returns the probability of the new point belonging to the individual clusters. The points were added to the cluster with the highest probability if it exceeded a threshold of 10%. Out of a total of 300,029 embeddings, 198,520 were classified as noise, while 10,685 embeddings were reassigned to clusters, following the described process. However, it became apparent that the application of a uniform static threshold across all clusters leads to the contamination of smaller as well as visually very complex clusters, such as cluster 162 (also see Sec. 3.1), which represents images in medical contexts. The cluster manifests an array of distinct scenarios, all semantically interlinked through the presence of medical themes. However, due to the heightened diversity of represented scenes, the cluster's demarcation from other topics becomes less pronounced and elevating the threshold leads to a reduction in its semantic distinctiveness. Please note that clusters that only occur in a single news video will not be displayed in the later graph visualization.

2.3. Bidirectional Graph Modeling

The next step toward a scalable visualization relates the clusters to each other by mapping them into a bidirectional graph structure. The nodes of the graph represent the previously generated clusters. For the creation of the edges, we experimented with two different approaches, resulting in two different graphs.

1. **Co-appearance frequency:** The first approach inserts an edge between clusters, if their constituent elements appear together in news videos and weights the edges based on the normalized amount of co-appearance. Frames belonging to a cluster must appear at least five times in a news video for the cluster to be counted as present in the video, and clusters must appear at least three times together for the edge to exist.
2. **Pearson correlation:** The second approach draws an edge if the clusters appear at least two times together and are weighted based on the Pearson correlation coefficient between the clusters. Edges with a weight under 0.3 were deleted.

The graph based on the number of co-appearances, aims towards an understanding of the

levels of centrality of motifs in the Tagesschau, so they can be experienced intuitively. Fig. 1 shows that the graph takes a circular form and concentrates around frequent topics or core motifs in central nodes. The periphery shows the less prominent clusters. A core topic could be the weather forecast and a core motif the intro sequence of each broadcast. The edges of this graph exhibit a high degree of intuitive comprehensibility. However, the mere co-appearance does not directly indicate a correlation between clusters.

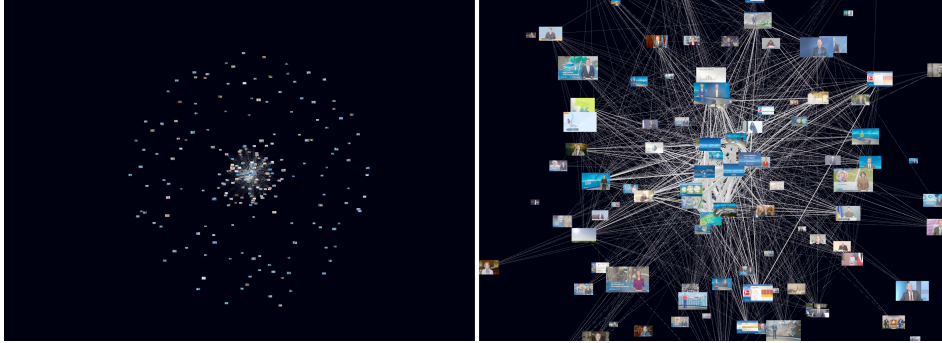


Figure 1: An overview on the left and a zoomed in view on the right of the co-appearance-based graph.

Therefore, the second graph aims towards the exploration of the correlation explicitly. In doing so, the visualization of this graph reduces the influence of the central topics of the Tagesschau. It creates a network that also highlights edges between clusters that are less frequent in the dataset, but still have high correlations. These clusters may then represent a specific interest, because they involve concrete issues and are less likely to be core motifs. This graph can be seen in Fig. 2.

Even though the two-fold graph structure is mainly used to create an interactive visualization, it also offers the possibility to use graph algorithms to further enrich the visualization with information that can be included as needed. Thus, communities were calculated with the *Clauset-Newman-Moore greedy modularity maximization*, including consideration of the weights. In this application, the community algorithm can be interpreted as a second clustering algorithm on the generated graph. In the later visualization, this will provide initial support for the identification of linked subgraphs. The communities, i.e. subgraphs, thus represent networks of frequently occurring motifs in daily broadcasts, or: networks of strong correlations. The resolution parameter, that if it is less than 1, favors larger communities and if it is greater than 1 smaller communities, was chosen to be 1.2.

2.4. Visualization and Interactivity

The created data structure is visualized in an interactive three-dimensional approach. The frontend of the web app is based on a Javascript framework built on ThreeJS/WebGL (<https://github.com/vasturiano/3d-force-graph>) and the force engine works with D3. The visualization is guided on a theoretical level by Arnold's and Tilton's third demand for distant viewing [2], with a pronounced emphasis on the *scalability* of the observational granularity.

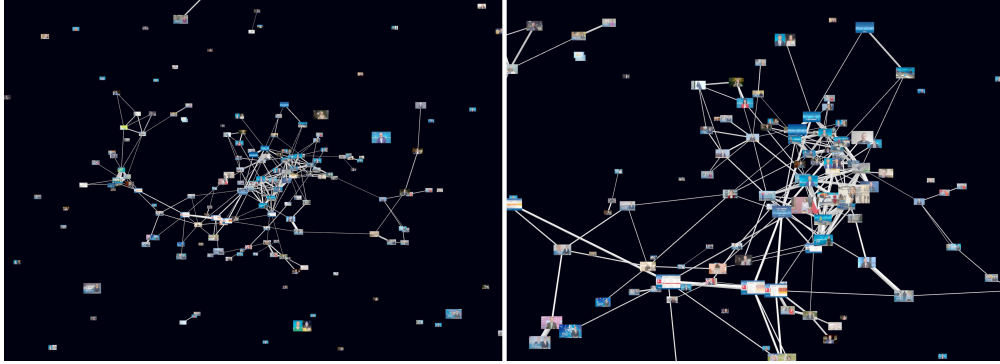


Figure 2: An overview on the left and a zoomed in view on the right of the correlation-based graph.

Additionally, it follows the principles by Shneiderman: “Overview first, zoom and filter, then details-on-demand” [19].

Zooming is possible via scrolling, allowing for an individually adjustable view magnification of the interrelationships. The cluster nodes are represented by a randomly selected frame of the cluster. The diameter of the edges depends on the edge weight and the graph structure represented in the 3-D space based on gravitational force influenced by the weight of the edge (Fig. 1 and Fig. 2). By clicking in the space and moving the cursor, the perspective of the camera can be changed, which allows the focus to be on different parts of the graph and, together with the zoom, provides a certain immersive ‘flight effect’. Individual nodes can be moved via drag and drop and the graph reorients itself depending on the force. By hovering over the node, the ID of the cluster can be explored. For a more detailed insight into the semantics of the cluster, a scrollable overlay can be opened by right-clicking on a cluster (Fig. 3), in which up to 80 random frames of the cluster are displayed. This feature enables a detailed examination and analysis of single clusters.

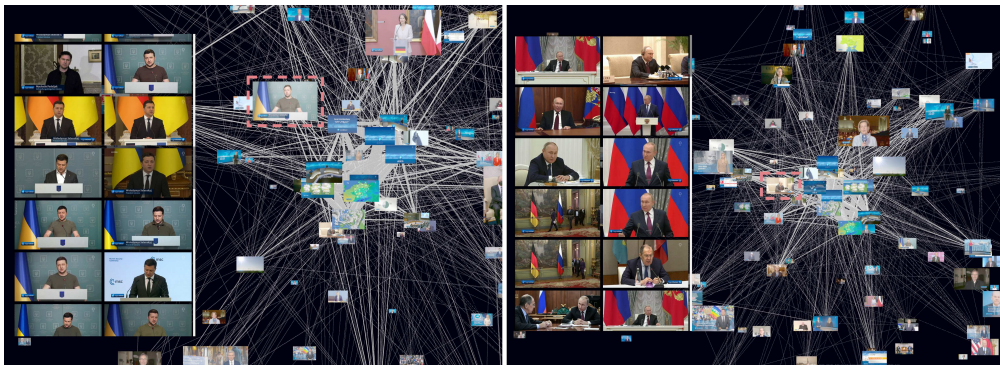


Figure 3: Illustrative instances of the scrollable overlay comprise frames associated with the selected cluster. The left side displays a cluster encompassing images of Volodymyr Zelenskyy derived from the correlation-based graph, while the right side showcases Putin and other Russian state officials sourced from the co-appearance-based graph. Clusters whose contents are displayed by right-clicking are framed with dashed lines.

To reduce the complexity of the cluster and to set a focus, it is possible to display only directly connected clusters by left-clicking on a cluster. Using the GUI, nodes can also be displayed that are connected via another edge. Another possibility to view sub graphs is the possibility to display one of the calculated communities in the graph via the GUI (Fig. 4).

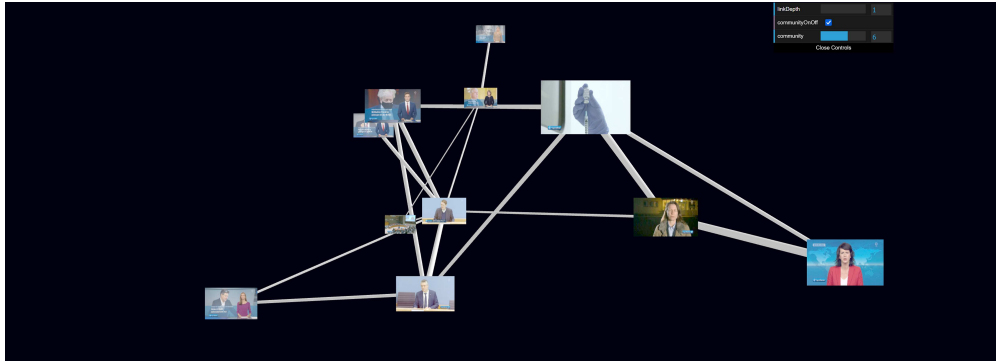


Figure 4: Filtering of the graph to only show community six of the correlation-based graph.

All of these features aim to create an interactive experience for the user that encourages continuous scaling of the viewing level, breaking down the barrier between micro- and macro-analytical viewing. Therefore, the developed visualization facilitates a comprehensive examination of large-scale patterns, pertaining to motifs and topics present in the data source, in this case specifically news videos. It serves as a crucial entry point for diving deeper into the detailed analysis of extensive volumes of video data, which would otherwise be impracticable to achieve manually.

3. Examples Analyses

The following examples are meant to demonstrate the exploratory analysis of video data with our methodology. Please note that all of these examples are using the correlation coefficient graph.

3.1. Topic: Covid-19

During the immersive exploration of the motifs of the Tagesschau in the visualization, cluster 162 stood out prominently. After scrolling through numerous images within the cluster and exploring their semantics via right-click, the majority of the images were confirmed to pertain to medical topics, predominantly related to Covid-19. Next, the focus was shifted toward clusters that showed direct linkage, which were accessed by left-clicking on cluster 162. As a result, four associated clusters emerged, each displaying distinct characteristics: the first (1) cluster depicted images from the *Bundespressekonferenz*, another (2) exhibited an anchor with a headline discussing Boris Johnson's *Partygate* scandal, there was an (3) interview cluster, and most notably, a cluster (4) presenting multiple instances of the *Meinung*-format, which translates to *opinion*. This exploration revealed a potential object of further analysis: the interaction

between topics related to Covid-19 and the utilization of a specific format that allows individuals to express subjective opinions. Fig. 5 shows the directly connected nodes of the medical cluster 162.

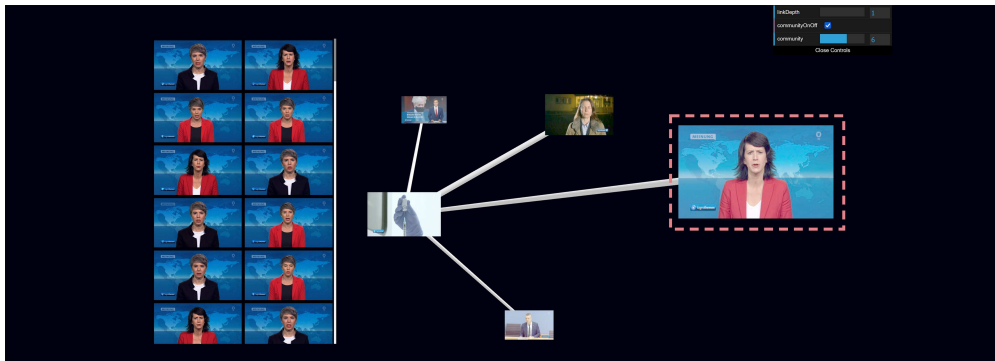


Figure 5: Directly connected clusters to the medical cluster 162. The overview shows the cluster containing frames from the “Meinung” format. Clusters whose contents are displayed by right-clicking are framed with dashed lines.

3.2. Topic: Ukraine War

The second example starts by analyzing the communities in the graph. Community 2 exhibits a sub-graph encompassing topics related to the war in Ukraine and another subgraph featuring a core element of the Tagesschau, namely, the weather forecast (Fig. 6).




Figure 6: Community two of the correlation graph; overview shows the cluster containing frames with maps of Ukraine. Clusters whose contents are displayed by right-clicking are framed with dashed lines.

Both sub-graphs are interconnected through a few edges, thereby forming a cohesive community together. Within the ‘Ukraine-sub-graph,’ three distinct clusters have been identified. The first cluster contains images of Volodymyr Zelenskyy, the current president of Ukraine. The second cluster consists of maps of Ukraine, and the last cluster showcases calls for donations to support Ukraine. At the periphery of the graph, two clusters featuring war reporters can be observed. Moreover, there is another link connecting to a cluster containing an an-

chorman and background reporting on war criminals, and another cluster showcasing Robert Habeck, the Federal Minister for Economic Affairs and Climate Action. The discovered connection of motifs appears to be more closely related to the contents of the Tagesschau. Hence, the visualization enables interactive exploration and presentation of complex themes that co-occur across multiple videos, allowing for the enrichment of these themes with additional information as required.

4. Discussion

The force simulation of the graph proved particularly beneficial in gaining an intuitive understanding of significant clusters. For instance, in the circular graph depicting co-appearance, it becomes evident that the force simulation places the core topics of the Tagesschau at the center. Indeed, prominent subjects within the dataset, such as the introductions, weather forecasts, charts displaying results of the German football league, and the previously described medical cluster 162, all converge at the center of the circular graph due to the force simulation. The analysis of the two graph variants should be considered in combination, as their complementary focus allows us to gain valuable insights into the motif network of the Tagesschau. Together, they offer a comprehensive perspective on the relationships and patterns within the dataset.

In the context of exploratory visualization, another crucial aspect to consider is that the ability to freely focus on specific elements while having abundant information might lead individuals to seek confirmation of their own hypotheses rather than openly discovering new connections between clusters. Similarly, the projection of semantics onto a cluster needs careful reflection and must be balanced and negotiated with one's own confirmation bias. Furthermore, the co-appearance in a news video as a unifying element and resulting correlations must always be considered in connection with certain basic conditions of the respective broadcast. This refers, for instance, to external factors that influence the program, such as unpredictable current events. However, this does not make the interaction of special topics, set focuses, motifs and special broadcast elements any less important. Finally, we want to highlight that during the creation of the graph it became evident that the parameters have a strong influence on the actual visualization. Their adjustment must therefore be critically examined to avoid the confirmation bias of a desired hypothesis.

5. Conclusion and Future Work

In this short paper, we present a novel approach that combines density-based clustering and graph modeling to create a scalable viewing application for the exploration of similarity patterns in news videos. Unlike most existing video analysis tools that focus on individual videos, our approach allows for an overview of a larger collection of videos, which can be further examined based on their connections or communities. By utilizing scalable viewing, specific subgraphs can be selected from the overview, and their respective clusters can be explored in more detail on the video frame level. The next step involves connecting the aforementioned

prototype for analyzing individual videos with the current scalable viewing prototype, enabling analyses of news videos at various levels of granularity.

Currently, the similarity function between frame clusters relies solely on CLIP embedding vectors. However, other information, such as written and spoken language, object and face recognition, color information, etc. can also be easily extracted from videos and will serve as additional information sources for future visualizations. While our current experiments focus on news videos, the scalable viewing approach can be applied to other video formats as well. Depending on the type and genre of the video, some clustering parameters may need to be adjusted since they are currently optimized for news videos.

6. Acknowledgements

This research was made possible through funding provided by the German *Federal Ministry of Education and Research* as part of the "FakeNarratives" project (16KIS1516).

References

- [1] T. Arnold and L. Tilton. "Distant viewing toolkit: A python package for the analysis of visual culture". In: *Journal of Open Source Software* 5.45 (2020), p. 1800.
- [2] T. Arnold and L. Tilton. "Distant viewing: analyzing large visual corpora". In: *Digital Scholarship in the Humanities* 34.Supplement_1 (2019), pp. i3–i16.
- [3] M. Burghardt, M. Kao, and N.-O. Walkowski. "Scalable MovieBarcodes—An Exploratory Interface for the Analysis of Movies". In: *IEEE VIS Workshop on Visualization for the Digital Humanities*. Vol. 2. 2018.
- [4] R. J. G. B. Campello, D. Moulavi, A. Zimek, and J. Sander. "Hierarchical Density Estimates for Data Clustering, Visualization, and Outlier Detection". In: *ACM Trans. Knowl. Discov. Data* 10.1 (2015).
- [5] G. Halter, R. Ballester-Ripoll, B. Flueckiger, and R. Pajarola. "VIAN: A visual annotation tool for film analysis". In: *Computer Graphics Forum*. Vol. 38. 3. Wiley Online Library. 2019, pp. 119–129.
- [6] S. Jänicke, G. Franzini, M. F. Cheema, and G. Scheuermann. "On Close and Distant Reading in Digital Humanities: A Survey and Future Challenges." In: *EuroVis (STARs) 2015* (2015), pp. 83–103.
- [7] S. Jänicke, G. Franzini, M. F. Cheema, and G. Scheuermann. "Visual text analysis in digital humanities". In: *Computer Graphics Forum*. Vol. 36. 6. Wiley Online Library. 2017, pp. 226–250.
- [8] D. Keim, G. Andrienko, J.-D. Fekete, C. Görg, J. Kohlhammer, and G. Melançon. *Visual analytics: Definition, process, and challenges*. Springer, 2008.
- [9] K. Kurzhals, M. John, F. Heimerl, P. Kuznecov, and D. Weiskopf. "Visual movie analytics". In: *IEEE Transactions on Multimedia* 18.11 (2016), pp. 2149–2160.

- [10] B. Liebl and M. Burghardt. “Designing a Prototype for Visual Exploration of Narrative Patterns in News Videos”. In: *INFORMATIK 2023 – Lecture Note in Informatics (LNI)*. Ed. by M. Klein, D. Krupka, C. Winter, and V. Wohlgemuth. 2023.
- [11] B. Liebl and M. Burghardt. “Zoetrope - Interactive Feature Exploration in News Videos”. In: *Annual International Conference of the Alliance of Digital Humanities Organizations, DH2022, Graz, Austria, July 10-14, 2023, Conference Abstracts*. Ed. by A. Baillot, T. Tasovac, W. Scholger, and G. Vogeler. 2023. DOI: 10.5281/zenodo.8107770. URL: <https://doi.org/10.5281/zenodo.8107770>.
- [12] L. Manovich. *Cultural analytics*. Mit Press, 2020.
- [13] L. McInnes, J. Healy, and S. Astels. “hdbscan: Hierarchical density based clustering”. In: *The Journal of Open Source Software* 2.11 (2017), p. 205.
- [14] F. Moretti. *Graphs, maps, trees: abstract models for a literary history*. Verso, 2005.
- [15] T. Ohm, M. C. Solà, A. Karjus, and M. Schich. *Collection Space Navigator: An Interactive Visualization Interface for Multidimensional Datasets*. 2023. arXiv: 2305.06809 [cs.CV].
- [16] J. Pause, M. Burghardt, and N.-O. Walkowski. “Scalable Viewing in den Filmwissenschaften”. In: *DHd 2019: Digital Humanities: multimedial & multimodal. Konferenzabstracts*. Ed. by P. Sahle. Frankfurt & Mainz, 2019, pp. 201–204.
- [17] K. Pustu-Iren, J. Sittel, R. Mauer, O. Bulgakowa, and R. Ewerth. “Automated Visual Content Analysis for Film Studies: Current Status and Challenges.” In: *DHQ: Digital Humanities Quarterly* 14.4 (2020).
- [18] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever. *Learning Transferable Visual Models From Natural Language Supervision*. 2021. arXiv: 2103.00020 [cs.CV].
- [19] B. Shneiderman. “The eyes have it: a task by data type taxonomy for information visualizations”. In: *Proceedings 1996 IEEE Symposium on Visual Languages*. 1996, pp. 336–343.
- [20] B. Shneiderman, C. Plaisant, M. S. Cohen, S. Jacobs, N. Elmqvist, and N. Diakopoulos. *Designing the user interface: strategies for effective human-computer interaction*. Pearson, 2016.
- [21] C. Tseng, B. Liebl, M. Burghardt, and J. Bateman. “FakeNarratives - First Forays in Understanding Narratives of Disinformation in Public and Alternative News Videos”. In: *9. Tagung des Verbands Digital Humanities im deutschsprachigen Raum, DHd 2023, Belval, Luxembourg and Trier, Germany, March 13 - 17, 2022*. Ed. by A. Busch and P. Trilcke. 2023, p. 138. DOI: 10.5281/zenodo.7715277. URL: <https://doi.org/10.5281/zenodo.7715277>.
- [22] C. Ware. *Information visualization: perception for design*. Morgan Kaufmann, 2019.
- [23] T. Weitin. “Scalable reading”. In: *Zeitschrift für literaturwissenschaft und linguistik* 47 (2017), pp. 1–6.