University of Zurich[UZH]

Gaudenz Halter, Alexandra Diehl, Renato Pajarola
and Barbara Flueckiger

Multi-level Visual Exploration of High-dimensional Spaces

TECHNICAL REPORT – No. IfI-2022.05

2022

University of Zurich
Department of Informatics (IFI)
Binzmühlestrasse 14, CH-8050 Zürich, Switzerland

Halter, Diehl, Flueckiger, Pajarola: Multi-level Visual Exploration of High-dimensional Spaces
Technical Report No. IfI-2022.05, 2022
Visualization and MultiMedia Lab
Department of Informatics (IFI)
University of Zurich
Binzmühlestrasse 14, CH-8050 Zürich, Switzerland
URL: http://vmml.ifi.uzh.ch/

Technical Report IfI-2022.05, Department of Informatics, University of Zurich 2022

Multi-level Visual Exploration of High-dimensional Spaces

Gaudenz Halter[†] Alexandra Diehl[†] Renato Pajarola[†] Barbara Flueckiger[‡]

Department of Informatics[†], Department of Film Studies[‡], University of Zurich, Switzerland

Figure 1: SankeyBridges visualization of *texture* (pink), *material* (blue) and *pattern* (yellow) features of LAWRENCE OF ARABIA (GBR / USA 1962, David Lean). The hierarchy of three feature groups are computed and juxtaposed in a radial layout. The relationships between different clusters of the feature groups can subsequently be revealed by means of a Sankey diagram inspired flow visualization.

Abstract

The increasing size and dimensionality of datasets in the humanities pose new challenges to scholars working with them, including establishing an overview over the dataset, connecting concepts, developing new hypotheses, and testing them. Material, pattern, and texture aesthetics in moving images is an attractive example of such multi-dimensional datasets in film studies, as an almost infinite number of combinations thereof are possible. Clustering techniques such as t-SNE are popular automated methods to organize these complex datasets, but they bring little or no-semantic meaning to their grouping strategies. We propose a novel interactive visualization technique for multi-level hierarchical exploration of clustered features, named Sankey-Bridges. *Our technique allows the users to (1) abstract local and global semantics from the automated methods, (2) extract relevant relationships, and (3) quantify them. Our technique is embedded in a system with other interactive visual components combined with exhaustive computational methods. The proposed solution is able to convey the global and local structure of high-dimensional clustered data sets and the relationship between different groups of features. The resulting visualization tool is embedded in the well-established VIAN [HBRFP19] research framework. We illustrate the benefits of our approach in the context of typical film researchers' investigation of relationships in high-dimensional spaces, and a wide range of qualitative analysis labels, with examples from an extensive film database.*

1. Introduction

Computational methods are an integral part of computer assisted data analysis procedures in digital humanities. What initially started with text processing evolved into a field entailing a large diversity in both, the methods used and their applications. In fact, computational methods have become a substantial advancement in recent digital humanities [EH13]. Consequently, large datasets have been collected and are continuously generated, which require efficient computer assisted as well as visual data analysis approaches to support interactive explorative hypotheses, their comparison and confirmation.

Using digital text, audio, image as well as video processing and analysis techniques, a large amount of feature vectors, auxiliary metadata or annotation data is commonly derived from the raw input to support subsequent use-guided classification, identification, comparison, visualization or search tasks. Therefore, these large and multivariate data sets pose many challenges to scholars, including how to establish an overview over the global structure of the data, how to detect reoccurring patterns in the data, how to derive hypotheses from observations, and finally how to test these hypotheses. Traditional methods to analyse a film corpus usually provide the scholar with a mental model of the results. An appropriate tool-set, grounded in theoretical and analytical concepts, extends these insights with new perspectives and hypotheses. Unlike other artifacts in the humanities, film is by definition of high dimensional nature, conveying both audio and visual information including motion, often embedding complex narrative structures, cultural context and stylistic idiosyncrasies of the artists involved in the production of a specific movie. These complex interactions in film aesthetics are particularly pronounced in the appearance of surface structures, such as textures, patterns and materials as arranged through art direction and costume design in the mise-en-scène in front of the camera [GH90, Gal11, NL12].

Therefore, textures, patterns, materials and how the surfaces interact with lighting are a crucial part of a film's aesthetics. Opposed to computer science, where generally the spatial value variation and patterns thereof describe a texture, following art history [Lan95] film studies differentiate between *textures*, being the three-dimensional spatial surface variation [LLZ15, ZJHJ16] and *patterns* being variations of color application on the surfaces of costumes, objects and environment. Although the two-dimensional optical representation of patterns and textures may look similar, they do not arise from the same qualities of a surface. The texture of a surface elicits *haptic* perception through cross-modality, and as such is often idiosyncratic of its material appearance [Flu20]. A surface's texture may be smooth or rough, it may be organized or may be random, geometrical, or ornamental. Due to how the light illuminates it, the texture is transformed into its two-dimensional representation on the image plane, which does however not exonerate it from its haptic qualities since it addresses the spectators' cross-modal perception. Patterns on the other hand are not connected to the haptic quality of the surface, their representation on the image plane roots purely in the effects from the variation of color application. As such, patterns can be described with similar terms as textures, they can be coarse or smooth, geometrical, floral, dotted, striped, checkered, or illustrative [Flu20]. As pointed out in [Flu20], the differentiation between texture and pattern is often difficult even for the human observer, and in many cases, they are not easily distinguishable. For example, painted stucco can exhibit both textural and pattern qualities.

In this paper we propose a new visualization, *SankeyBridges*, to convey the relationship between different feature vectors by portraying their hierarchy. Our visualization allows researchers to develop insights into relationships in their data set and to test hypotheses. Our novel visualization technique is embedded in an interactive visualization tool that allows the user to visually explore multi-dimensional data sets and obtain new insights into relationships and structures to be found in the data set. We further exemplify our approach by studying a number of use cases typical in film research where we focus on the above outlined problem of texture, pattern and material in film.

The main contributions of our paper are:

- A new visualization method *SankeyBridges*, which displays hierarchically organized t-SNE embeddings in a circular layout and allows the user to identify relations between different feature groups within a data set by means of a Sankey-diagram inspired flow glyph.
- A fully working prototype which leverages said visualization to allow film scholars to identify relationships between *material*, *pattern* and *texture* feature probabilities.
- An elaboration on several use cases designed in conjunction with film researchers that showcase the potential of our technique.

2. Related Work

2.1. High-Dimensional Data Visualization

Dimensional reduction is a key task when dealing with high-dimensional data sets in visualization. Since our proposed solution organizes features into groups and depicts their relationships in a in a two-dimensional visualization, we reduced the dimensionality of the input data set accordingly. Many methods have been proposed for this task, such as Principal Component Analysis (PCA), MultiDimensional Scaling (MDS), Self-Organizing Maps (SOM), Uniform Manifold Approximation and Projection (UMAP), and t-distributed Stochastic Neighbor embedding(t-SNE) and its variants. Today, t-SNE is used in a wide variety of fields handling about high-dimensional data, such as analyzing single nucleotide polymorphisms (SNPs) [Pla13], time dependent data [RFT16] or even in combination of graph layouts [KRM*17]. A variant of t-SNE, BH-SNE [VDM14], has been applied to high-dimensional data sets of various fields, such as single-cell analysis [ADT*13], and spatial gene expression organization in the brain [MvdGvdM*15], and Q-SNE was developed for dimensionality reduction of large document data [IM15]. A hierarchical SNE has been proposed in [PHL*16] to allow a global-to-local navigation through the data set by means of hierarchically structured embeddings, allowing for interactive browsing and mitigating the high computational costs that SNE algorithms typically consume. An approach which has already been successfully applied to cytometry data [vUHP*17]. Recently, tree-SNE has been introduced [RPH20], which stacks one-dimensional t-SNE embeddings on top of each other, revealing hierarchical structures within the data. We use the Barnes-Hut approximation of the t-SNE *BH-SNE* to minimize the computation time between a query of the user and returning the computed hierarchy. Our proposed solution picks up this idea of hierarchically organizing several two-dimensional t-SNEs, and therefore allowing a global-to-local visualization approach.

Other approaches modeled paths as clustered high dimensional data sets and mapped them using reduction techniques such as t-SNE and UMAP to visualize trajectories and reveal hidden path patterns [HSS*20]. We share the use of visualizations through clustered relations as a key element, in our approach we exploit the hierarchical organization of our clustering method, BH-SNE. Furthermore, other radial visualizations such as RadViz [HGP99] are very popular visualization techniques to visualize high-dimensional data sets. For example recently, RadViz++ has been proposed as an improved method [PT19], that allows to interactively aggregate, separate and filter variables, as well as see the impact of the layout in real time. In our approach we also use a circle layout to show summaries of the high-dimensional problem. However, we do not place the variables along the circular design space. For our design we chose a star-layout for the three main dimensions (pattern, texture, and material) to maintain the inherent hierarchical structure of our BH-SNE method in between dimensions.

2.2. Hierarchical Data Visualization

Hierarchical data structures and representations have largely been studied in areas such as graph visualization [VLKS*11, VBW15], hierarchical tree structures [LZD*19, RPH20, SHS10], network visualization [HLT*20], or machine learning [KK96, HVP*19, CMJ*20]. In particular, hierarchical edge bundling technique are suitable for the visualization of adjacency relations in hierarchical data [Hol06]. Our hierarchical edge bundling technique is inspired by this, but we expanded it by adding the Sankey diagram metaphor to the bundles. Moreover, we also integrate a radial visualization layout, similar to chord diagrams, combined with the Sankey diagram concept resulting in our *SankeyBridges*, that show and quantify the connection among different groups, and enhanced it by focus+context interactions. Interactive focus+context visualization techniques have been proposed for the exploration of HSNE clusters [HVP*19]. Furthermore, PLANET is a radial layout algorithm for network visualization [HLT*20] based on previous work from Pavlo [PHS06]. In our approach, we also use parent-centric radial layout of spanning trees as a method to overview different levels of our clustering hierarchies. The main difference of our work is that we encode the relations among clusters and the strengths of these relations using Sankey-like diagrams.

2.3. Texture, Pattern and Material Analysis

In (digital) film studies, the (computational) analysis of the visual content of movies, shot segments or individual frames, is a fundamental task. With the natural affinity to visual representations, the research community has in particular been striving to establish ways to coherently visualize the content of movies and ensembles thereof. An extensive list of current methods can be found in the literature, including an investigation of the usability for film scholars [Stu16], as well as reflections on theoretical, epistemological and historical concepts in digital methods for film studies [Ole17, Hef18, Hef16]. In recent years a growing amount of work has been dedicated to depicting color distribution in films, mostly from a distant reading perspective [Mor13]. This includes, for example, the use of radial palettes for color visualizations [Bro11] or the MovieBarcodes [BKW16, BHE*17] for further analyses of dialogues in an innovative way. Normalized mean images along the temporal axis have been proposed to visualize the spatial color content of the complete movie [Fer17], and an interactive timeline visualization to allow comparisons between dialogue and color schemes of a movie [HSSS17]. In particular, the movie annotation system VIAN [HBRFP19], with its focus on colorimetry, as well as previous work on the combination of clustering and interactive visual analysis of color information of single movies and corpus visualizations [Flu17] are loosely related to our work [FH20].

Another group of related work is focusing on techniques for exploratory visualization. Exploratory visualization emerged to assist scholars in hypothesizing and speculating about relationships hidden in increasing dimensionality and size of data sets. In humanities such techniques have become popular as well [HFM15] with many applications such as tag clouds [VWF09], investigating relationships between text documents [SGL08], exploring news articles [DCCW08] and many more. What most of these systems in the realm of information visualization have in common, is that a visualization is linked to a database allowing the user to view a selection in detail, either by employing a multi-view layout [Rob05] or providing some query-driven access to the data set [SGL08].

3. Visual Design

The *SankeyBridges* visualization aims to convey relationships between different feature vectors of data sets by leveraging their hierarchical structure. In this section, we will elaborate in detail how it is created. Fundamentally, *SankeyBridges* is a combination of two main components: A radial and hierarchically organized layout of low dimensional embeddings of the feature spaces, and interactive flow curves which connect these embedding based on their shared items as illustrated in Fig. 1.

3.1. Computing the Hierarchy

Given a data set containing three feature groups $\boldsymbol{F}_1, \boldsymbol{F}_2, \boldsymbol{F}_3$, e.g. as in our use case, the output class probabilities of three convolutional networks, a subsequent step is to present both the global and local structure of the data set by means of dimensionality reduction in a single view. This single view approach is important for two reasons, on one hand it allows the observer to gain information about the groups feature diversity without navigating, on the other hand it allows us to depict relationships between different feature groups using a Sankey-like flow visualization. We achieve this hierarchy structure by employing a recursive function which builds a hierarchy of t-SNE embeddings.

For each feature group $\boldsymbol{F}_i$, the hierarchy is constructed by first computing a low-dimensional t-SNE embedding of the complete data set followed by grouping the data points of this embedding into k different regions, and recursively repeating these steps for each of the resulting clusters until the desired depth d is reached.

To ensure the formation of a low number of sizeable groups within each embedding, we set the t-SNE perplexity to be relatively large, specifically, we compute the perplexity by the number of items in the embedding with $p_d = \sqrt{N_d}$ where N_d is the number of items within the group at depth d. This approach has already shown to be a good estimate [RPH20]. The generated output is a tree of t-SNE embeddings which transition from global to local structures with increasing depth of the tree.

To clearly illustrate this hierarchical relationship, we employ a circular layout, which is first divided into the different feature groups, and which shows multiple layers radially outwards corresponding to the depth of the hierarchy. In Fig. 2, the clustering of the first three layers of the feature group *Material* is shown. The central layer contains a root t-SNE embedding incorporating the complete data set, each consecutive layer contains the k child embeddings which are the result of the clustering performed on their parent layer. Therefore, the more the layer depth increases, the more the individual embeddings differentiate into coherent groups of a certain feature probability combination.

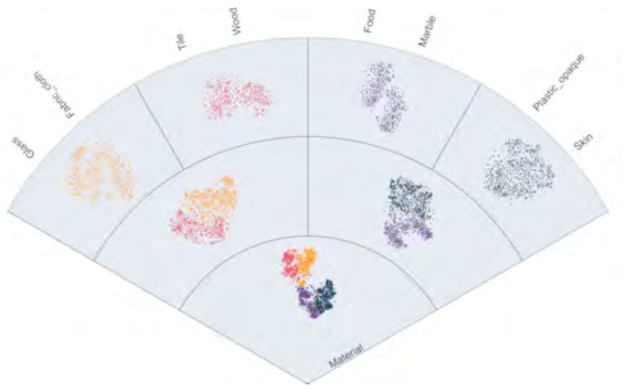

Figure 2: Hierarchical visualization of the *Material* feature group. In the central area of the circle the complete data set is embedded into a two-dimensional space based on the material probabilities. With each subsequent layer, the data points are divided by means of a clustering algorithm and embedded anew. The color indicates the final clustering of the data points as associated to the leaves nodes of the tree.

3.2. SankeyBridges

Routinely, the user is not only interested in a single feature group and how the data set expresses its set of features, but in how the expression of different feature groups are related. For example, a user may want to investigate if certain materials in set and costume design are related to certain patterns or texture combinations. Additionally, in the case of digital humanities approaches, items are often fortified with large and extensive keyword annotations or even semantic tags. This characteristic raises questions such as how these verbal keywords correlate to the expressed visual features yielded by the various models.

We depict this relationships between different feature groups by means of drawing Sankey-like flow visualizations, *SankeyBridges*, from the selection done by the user to every embedding that contains data points of the selection. The rationale behind this method builds on the hypothesis that the relationships between the different feature groups can be identified by looking at the distribution of feature expression across the hierarchy of them. For example, given a set of screenshots extracted from a movie set in an old castle, we would expect to see that screenshots containing the material *stone* would often express the texture *cracked*, since this is a typical texture found on castle walls. In such typical cases, the screenshots would not be evenly distributed throughout the texture and material hierarchy, but instead be highly concentrated in the respective sub regions which contain stone and cracked texture. If the movie also contains a second scene which takes place in a mountain valley, we could expect to find a second texture often associated to the material *stone* for example *bumpy*, since rock appearing in nature exhibits coarse, bumpy surfaces, thus with increasing combination diversity of the material rock with textures, the SankeyBridges in the texture hierarchy would become more evenly distributed. Fig. 4 illustrates this with two different selections in the *pattern* feature group.

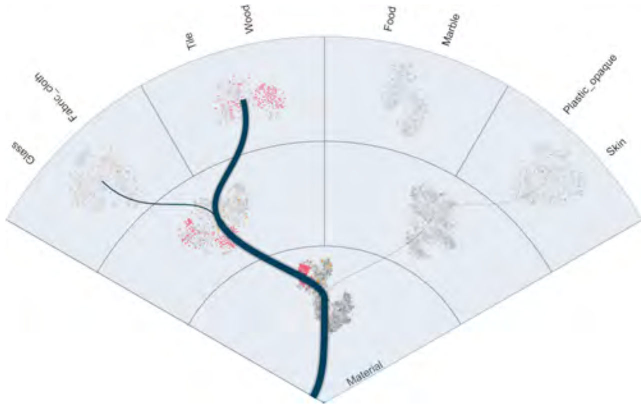

Figure 3: An incoming SankeyBridge flow visualization in one feature groups. The width of the flow indicates the number of shared items between the source and destination cluster.

Our visualization technique highlights these relevant relationships by drawing SankeyBridges through the hierarchy tree bridging the different feature groups. The thickness of a each SankeyBridge flow line represents the strength of the relationships defined by the number of items the destination cluster shares with the initial selection as shown in Fig. 3. Given a selection in the $\boldsymbol{F}_1$ feature group, if the relationships distribute evenly throughout the $\boldsymbol{F}_2$ feature hierarchy, shown in Fig. 4a, one can assume that there is no combination of textures which is correlated to the selected pattern probabilities. In contrast, if there exists only few strong SankeyBridges from the selection to a destination cluster which is a leaf of the hierarchy, as shown in Fig. 4b, an association can be assumed between the distributions of $\boldsymbol{F}_1$ and $\boldsymbol{F}_2$. Such relationship queries can be performed either by using the lasso tool or by clicking on the area where a certain cluster is located.

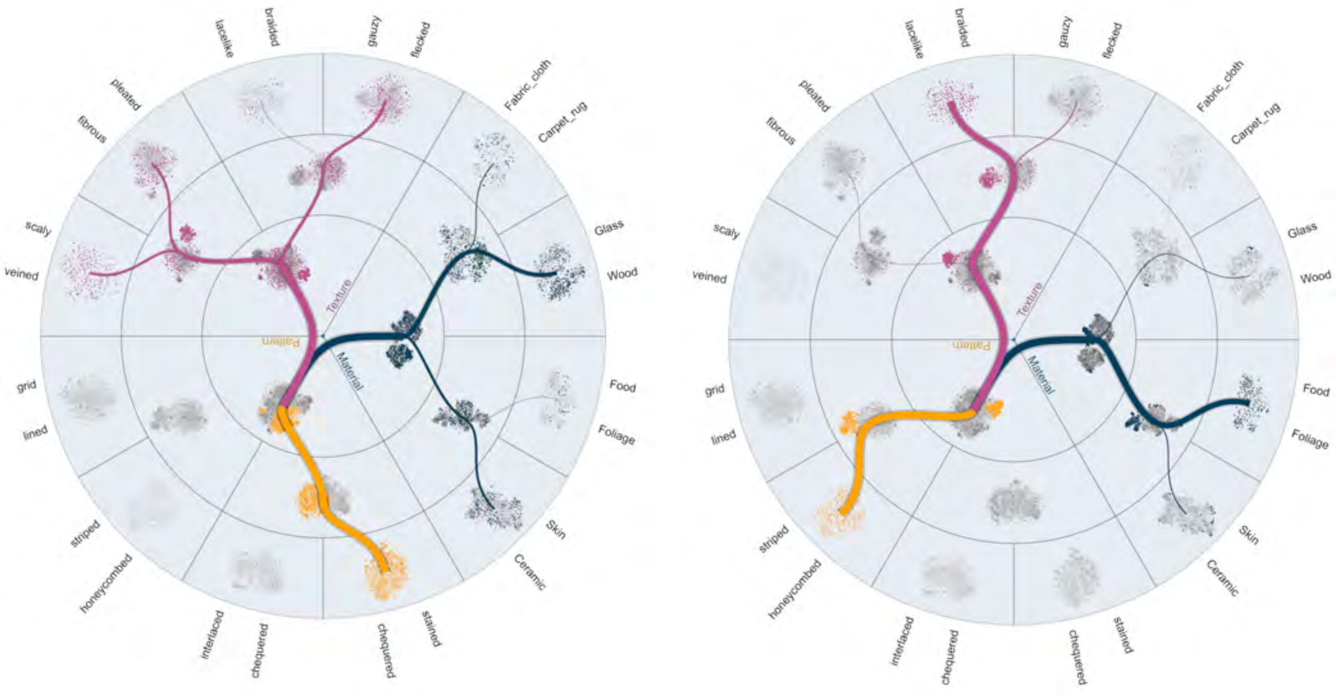

(a) Weak relationship

(b) Strong relationship

Figure 4: Comparison between a weak and a strong relationship indicated using SankeyBridges with three feature groups (red, green and blue). In (a) a weak relationship is conveyed by an even division of the Sankey flow from the selected region (yellow) to other feature groups in the hierarchy. In contrast (b) shows a skewed distribution to only two other feature groups, which indicates a strong relationship.

4. Prototype Application

In this section, we will present our prototype application which applies the *SankeyBridges* visualization method on a large film data set curated, segmented and annotated by domain experts. Using multiple additional widgets we enable the film scholars to explore based on *texture*, *pattern* and *material* and cross-reference it with their database.

4.1. Requirements

As noted in the introduction, it is very common in humanities to parse the information at hand into smaller, coherent fragments with respect to the specific research question [Hah09, CBC12]. In the case of the data set at hand, several types of such fragments exist on the micro, meso and macro level for screenshots, temporal segments (shots, sequences), individual films and the film corpus or subcorpus. Temporal segments were parsed with reference to a coherent color distribution in a sequence of shots. Finally selected frames were picked for each segment to represent the aesthetic characteristics in the selected shot, such as image composition, camera angle, lighting, staging,. In VIAN [HBRFP19], temporal segments are then enriched by *coding* them with externally defined ontologies or vocabularies.

Since our system embeds into this research environment, it has been an important requirement during the development to have direct access from our visualization to the corresponding items of these three levels. In Figure. 5, we illustrate how our application embeds into the existing research infrastructure. We provide two components to the system: On the feature acquisition side, an automated feature extraction has been implemented which extracts the predictions for the different materials, textures and patterns. On the visualization side we developed a functional visualization application.

4.2. Image Features

In film studies, it makes sense to split the spatial information of an image into color variations on surfaces referred to as *patterns* and three-dimensional and thus haptic variations of surfaces, referred to as *textures* respectively. Profilmic materials in set and costume design are significant elements of a film's style related to the period of its production, the technical color film process applied and the cultural context in architecture, design and fashion. To investigate patterns, textures and materials in film, we trained three convolutional neural networks (CNN) based on the InceptionV3 architecture [SVI*16] to predict the pattern, texture and material features of a given image. In case of material we used the OpenSurfaces [BUSB13] data set as training data. For pattern and texture, we took the classes from the DescribableTextures-Dataset [CMK*14] and split them into the two categories *textures* and *patterns* based on their visual characteristics. Subsequently, we used transfer learning with a base model pre-trained on the imageNet data set [GRM*18] and retrained it using the previously de-

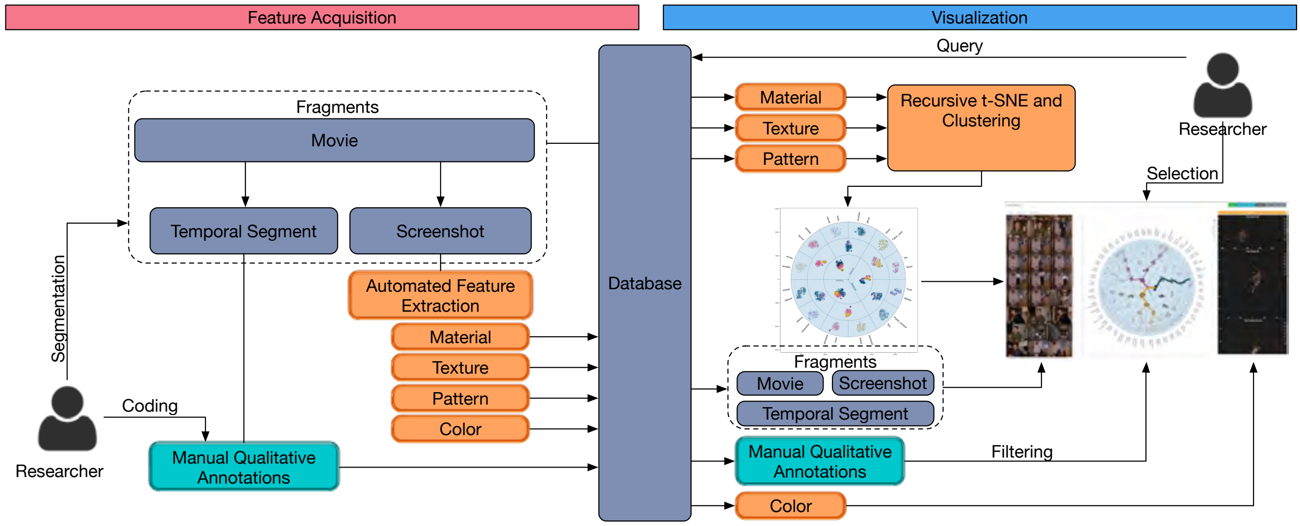

Figure 5: Overview over the pipeline of the film analysis and data flow in our visualization. The workflow consists of a feature acquisition stage, where a single movie item is selected, segmented into fragments (gray) and screenshots. This stage is followed by a manual, qualitative classification of the individual temporal segments based on a large controlled vocabulary (turquoise), while the screenshots are then used to automatically extract feature vectors. During the visualization stage, the material, texture and pattern feature vectors are used to create a tree of t-SNE embeddings based on a user's query. Using the selection tool, the user is able to generate SankeyBridges between the different feature groups and further investigate it by cross-referencing the underlying fragments using the item inspection widgets as well as the color distribution visualizations.

scribed training data and additional fine tuning data manually annotated. After training, we applied these three classification models on a data set of 170'000 screenshots from mostly feature films throughout film history from 1895 to 1995, and kept the class probabilities as feature vector for later visualization.

4.3. Visualization Application

To test our solution, we implemented a fully working prototype visualization application targeting film scholars. Although our *SankeyBridges* visualization can yield a lot of insights without additional visual components, it is crucial to embed such systems into the existing research framework of the domain experts who test the application, such that they can cross-reference the patterns they observe using the prototype directly to their database and tools they are already familiar with.

Our prototype, shown in Fig. 6, consists of an interactive central SankeyBridges visualization which displays all items returned by a query as hierarchically organized t-SNE embeddings as described previously. Such a query may encompass an arbitrary set of screenshots stored in the scholars database we described more in details in Sec. 5, which may either be entered by a text box in the search bar, or by means of a exported query JSON string from their database platform.

Since the output of our models are probabilities for different material, texture and pattern classes, and because these classes are human interpretable, such as *wood* or *striped*, extended our visualization: For each leaf embedding, we extract the two classes which exceed the mean of the visualized items the most, and place them at the border of the data set.

In addition, we added several views which display details about the current selection. On the left side, the selected items are shown on the three different conceptual levels the domain experts work with: In a first tab, the selected screenshots are directly displayed, a second tabs displays the temporal segments within a movie the screenshots belong to, and finally, the respective projects are displayed. All of these widgets serve as a direct link to the web platform of the film scholars, which hosts the data for the application.

On the right side, additional colorimetric information is given about the selection, namely the color distribution in respect to the color space and the production year of the selected movies.

4.4. Implementation

The data preprocessing, feature vector extraction and backend of the final visualization have been implemented in Python3. The CNNs have been implemented using Keras in conjunction with TensorFlow, and OpenCV for the image processing tasks. For the clustering tasks, distance metrics, and statistics, scikit-learn has been used. Finally, Flask is used for the server side of the visualization in combination with SQLAlchemy for the database binding.

The visualization itself is implemented in JavaScript. BokehJS has been used for the SankeyBridges and color distribution visualizations. The user interface used JQuery in combination with bootstrap as a layout framework.

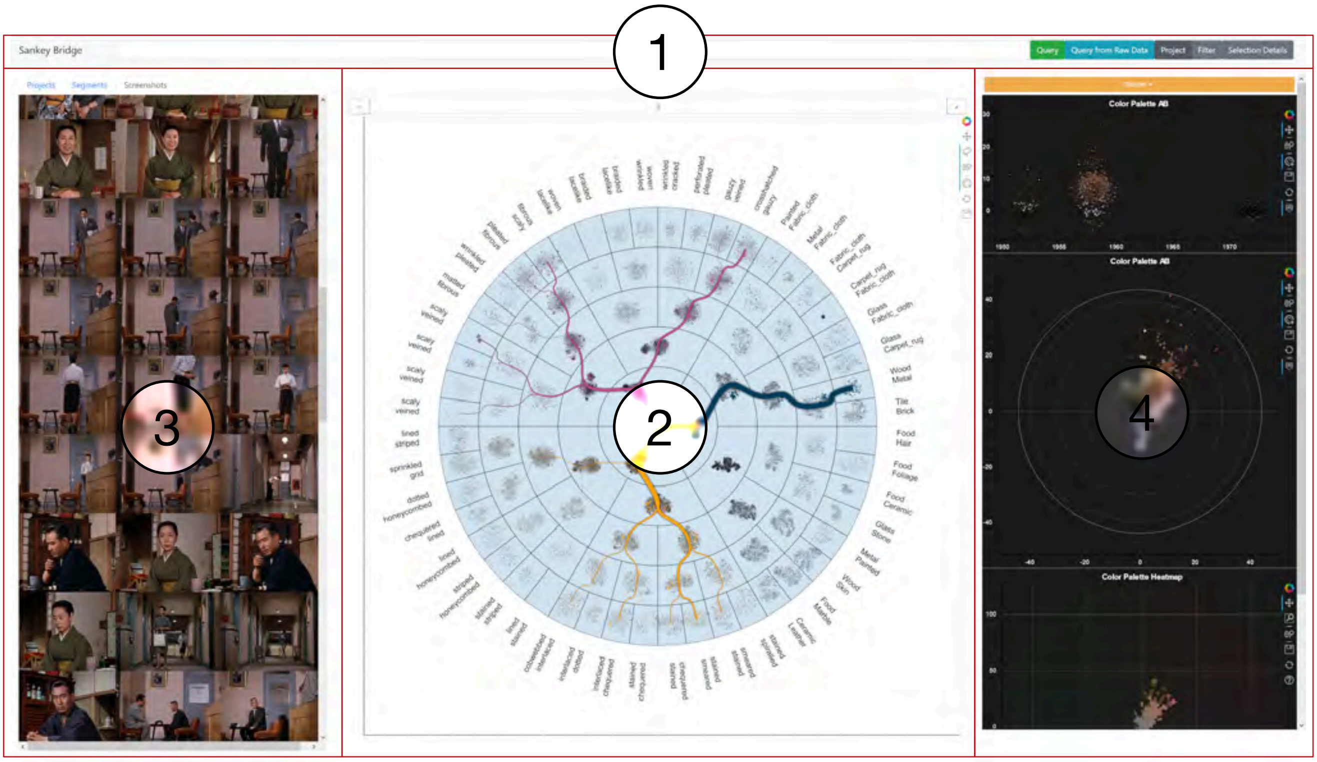

Figure 6: The main view of the implemented solution. The search bar (1) allows the user to create a new visualization based on certain movies and metadata, as well as filtering the visualization based on qualitative tags. The main visualization (2) gives shows the current hierarchical layout of t-SNE embeddings and gives the user several possibilities to interact with it, such as lasso selection or zooming. Once a lasso selection has been performed, the details panels (3) show the screenshots, temporal segments and movies which belong to the selection. An additional color information widget (4) gives information about the selections' color distribution in terms of color space and production year.

5. Use Cases

Two use cases of our tool illustrate the usability and results generated in *SankeyBridges* for visualizations in film studies analysis and research. They focus on two sets of approaches.

5.1. Paris, Texas

As a first example, we computed our visualization for PARIS, TEXAS, (FRA 1984, Wim Wenders). The film shows a distinctive style that results from the collaboration between director Wim Wenders and cinematographer Robby Müller who draw on American iconography established in photography. Three main characteristics can be retrieved by looking at the hierarchy of the of the materials feature group embeddings, shown in Fig. 7:

1. rough textures in rocky landscapes and grassy steppes are associated to the male protagonist in the first third of the film. They evoke the photographic style of Walker Evans and Dorothea Lange for the Farm Security Administration in the 1930s.
2. cozy interiors with floral patterns on wallpapers and the female character who provides warmth and a safe space to the protagonist
3. colored illumination in smoothly textured, non-patterned interiors that create a pastiche of William Eggleston's color photography of the 1970s and 1980s.

In Fig. 7 the computed hierarchy of the movie screenshot is shown. By inspecting the hierarchy using the lasso tool, we can see, that the grouping is sound with the three main locations the movie is set. Looking at the SankeyBridges between the selected outdoor area (Fig. 7 (3)) we can further see that the main flow leads to the cluster (5) which includes a large super set of screenshots with generally coarse textures.

5.2. Lawrence of Arabia

LAWRENCE OF ARABIA (GBR / USA 1962, David Lean) is a classic widescreen film of the 1960s that depicts the struggle and decay of British writer T. E. Lawrence in the Arabian desert. Textures and image compositions are associated to Lawrence's mental state.

The tool reliably extracts different sets of image compositions, textures, surfaces and patterns, shown in Fig. 8.

1. grandiose, wide panoramas of the sandy desert with blue skies or sunsets that are characteristic for the protagonist's sense of supremacy and invincibility at the beginning of the film. They

Figure 7: The resulting SankeyBridges visualization of the movie PARIS, TEXAS (FRA 1984, Wim Wenders). (1) Rough textures in rocky landscapes and grassy steppes, (2) cozy interiors providing warmth for the protagonist, (3) colored illumination in smoothly textured, non-patterned interiors. The smoothness is of the texture is also indicated by the flow diagram leading to the region marked with gauzy. The group in (1) further splits into pure landscapes (4) and frames containing grass and characters (5).

depict a European's romantic idea of the desert with smooth textures and flowing lines.

2. tumultuous war scenes with thousands of Arabian warriors on horses that create atmospheric diffusion through swirling dust and sand. On the level of image compositions these scenes combine coarser textures with dotted sprinkles, all in earthy, largely achromatic tones.
3. the desert as a hostile environment with coarsely textured, hard surfaces express the increasing and life threatening conditions that take their toll on the protagonist's physical and mental health.
4. one interesting feature is visible in a collection retrieved with the label *banded*: it depicts Lawrence's appropriation of Arabic clothing with the *keffiyeh* and a golden *agal* wrapped around his head. While this costume is iconic for T.E. Lawrence's overblown sense of heroism, it bears an increasingly ironic undertone when it is covered in sand and dust with Lawrence's worn-out face

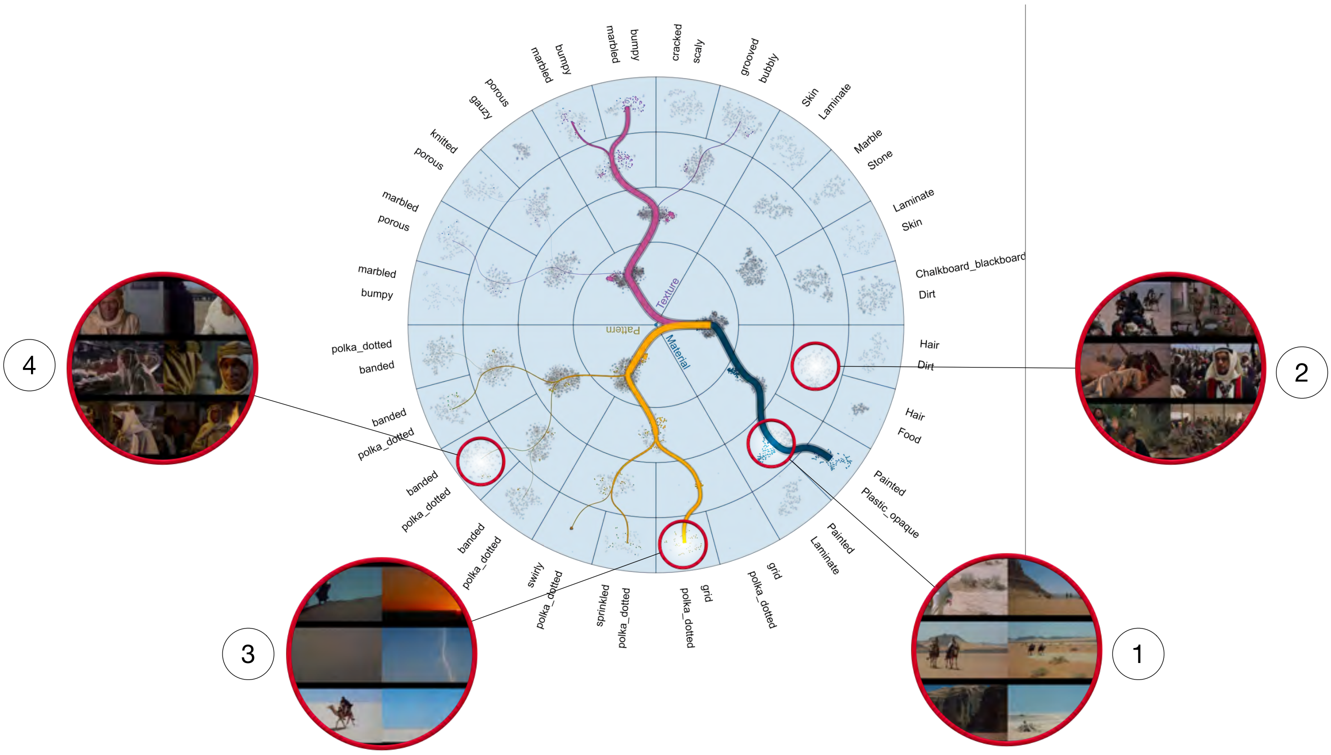

Figure 8: SankeyBridges visualization of LAWRENCE OF ARABIA (GBR / USA 1962, David Lean). (1) Wide panoramas of the sandy desert with blue skies, (2) tumultous war scenes with thousands of Arabian warriors, (3) the desert as hostile environment and (4) Lawrence's *keffiyeh* and a golden *agal* wrapped around this head. The current selection lies on the pattern cluster (3), and the Sankey flow visualization indicates a strong relationships to few clusters in material and texture. Since the desert the screenshots depict the empty and hot desert, only few materials and different textures appear.

6. Discussion

Relationships between different feature groups are an interesting target for visualization applications, especially in a data set where the different groups are known to interact. *Texture*, *material* and *pattern* in moving image are an excellent example for such feature groups, as their composition is not highly dependent. The texture of a surface is influenced by the material, and certain patterns can occur in combination with certain textures.

In this paper, we demonstrated how SankeyBridges proves to be an effective tool to investigate multi-dimensional data sets in the context of film analysis. It has been shown with the example of perceptual texture, material and pattern classification that a global-to-local navigation of the data set using recursively computed t-SNEs helps the scholar to identify trends in the global as well as more nuanced changes in the local structure of the data set regarding a single feature group. Using a Sankey diagram inspired flow visualization, we then bring several such hierarchies into relationship to each other and therefore allow the researcher to investigate their interactions by means of interactive selection. Our prototype targeting film scholars embeds into the existing research infrastructure enabling them to further investigate their hypotheses in their extended tool-set.

7. Limitations and Future Work

Our prototype exhibits some limitations which we plan to address in the future. First, we have not integrated the full query system of the researchers' platform. Instead, we allow the user to perform a query on the platform and export the retrieved items in a JSON file to pass to our prototype. Second, although we have shown that our CNN models are able predict the classes of the images well enough to form meaningful groups in regards to *texture*, *pattern* and *material* to create a hierarchy, there is room for improvement. Specifically, we did not yet source our texture and pattern classes in a catalog approved by film studies, but by the data set available, hence, the relevance of several classes for this application is disputable while some relevant ones are missing. This is something we plan to implement in the future. However, SankeyBridges is not limited to the realm of film aesthetics but can be used in other domains for the discovery of relationships, thus, we aim to apply our method to other data sets in the future.

Acknowledgements

This project was partially supported by a Swiss National Science Foundation (SNSF) Sinergia research grant (project no. CRSII5_198632).

References

[ADT*13] AMIR E.-A. D., DAVIS K. L., TADMOR M. D., SIMONDS E. F., LEVINE J. H., BENDALL S. C., SHENFELD D. K., KRISHNASWAMY S., NOLAN G. P., PE'ER D.: Visne enables visualization of high dimensional single-cell data and reveals phenotypic heterogeneity of leukemia. *Nature Biotechnology 31*, 6 (2013), 545–552. doi:10.1038/nbt.2594.

[BHE*17] BURGHARDT M., HAFNER K., EDEL L., KENAAN S.-L., WOLFF C.: An Information System for the Analysis of Color Distributions in MovieBarcodes. In *Proceedings International Symposium on Information Science*, vol. 70. Verlag Werner Hülsbusch, Glückstadt, 2017, pp. 356–358.

[BKW16] BURGHARDT M., KAO M., WOLFF C.: Beyond Shot Lengths. Using Language Data and Color Information as Additional Parameters for Quantitative Movie Analysis. In *Digital Humanities 2016. Conference Abstracts* (Kraków, 2016), Jagiellonian University & Pedagogical University, pp. 753–755.

[Bro11] BRODBECK F.: Cinemetrics. Film Data Visualization, 2011. OS. URL: http://cinemetrics.fredericbrodbeck.de/.

[BUSB13] BELL S., UPCHURCH P., SNAVELY N., BALA K.: OpenSurfaces: A richly annotated catalog of surface appearance. *ACM Trans. on Graphics (SIGGRAPH) 32*, 4 (2013). doi:10.1145/2461912.2462002.

[CBC12] CUTTING J. E., BRUNICK K. L., CANDAN A.: Perceiving Event Dynamics and Parsing Hollywood Films. *Journal of Experimental Psychology Advance online publication* (2012). doi:10.1037/a0027737.

[CMJ*20] CHATZIMPARMPAS A., MARTINS R. M., JUSUFI I., KUCHER K., ROSSI F., KERREN A.: The state of the art in enhancing trust in machine learning models with the use of visualizations. In *Computer Graphics Forum* (2020).

[CMK*14] CIMPOI M., MAJI S., KOKKINOS I., MOHAMED S., VEDALDI A.: Describing textures in the wild. In *Proceedings IEEE Computer Vision and Pattern Recognition* (2014).

[DCCW08] DÖRK M., CARPENDALE S., COLLINS C., WILLIAMSON C.: Visgets: Coordinated visualizations for web-based information exploration and discovery. *IEEE Transactions on Visualization and Computer Graphics 14*, 6 (2008), 1205–1212.

[EH13] ELL P. S., HUGHES L. M.: E-infrastructure in the humanities. *International Journal of Humanities and Arts Computing 7*, 1-2 (2013), 24–40. doi:10.3366/ijhac.2013.0079.

[Fer17] FERGUSON K. L.: Digital surrealism: Visualizing walt disney animation studios.

[FH20] FLUECKIGER B., HALTER G.: Methods and Advanced Tools for the Analysis of Film Colors in Digital Humanities. *Digital Humanities Quarterly 14,4*, Special Issue Digital Humanities & Film Studies. Analyzing the Modalities of Moving Images (2020).

[Flu17] FLUECKIGER B.: A digital humanities approach to film colors. *Moving Image: The Journal of the Association of Moving Image Archivists 17*, 2 (2017), 71–94. doi:10.5749/movingimage.17.2.0071.

[Flu20] FLUECKIGER B.: Textures, Patterns and Surfaces in Color Films. *Colour Turn 2,1* (2020).

[Gal11] GALT R.: *Pretty. Film and the Decorative Image*. Columbia University Press, New York, May 2011.

[GH90] GAINES J., HERZOG C.: *Fabrications. Costume and the Female Body*. Routledge, 1990.

[GRM*18] GEIRHOS R., RUBISCH P., MICHAELIS C., BETHGE M., WICHMANN F. A., BRENDEL W.: Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. *arXiv preprint arXiv:1811.12231* (2018). doi:arXiv:1811.12231.

[Hah09] HAHN S.: Filmprotokoll Revisited. Ground Truth in Digital Formalism. In *Digital Formalism: Die kalkulierten Bilder des Dziga Vertov*, vol. 55,3. Böhlau Verlag, 2009, pp. 129–136.

[HBRFP19] HALTER G., BALLESTER-RIPOLL R., FLÜCKIGER B., PAJAROLA R.: VIAN: A visual annotation tool for film analysis. *Computer Graphics Forum 38*, 3 (June 2019), 119–129. doi:https://doi.org/10.1111/cgf.13676.

[Hef16] HEFTBERGER A.: *Kollision der Kader. Dziga Vertovs Filme, Die Visualisierung ihrer Strukturen und die Digital Humanities*. No. Band 2 in Film-Erbe. Edition Text + Kritik, München, 2016.

[Hef18] HEFTBERGER A.: *Digital Humanities and Film Studies. Visualising Dziga Vertov's Work*. Quantitative Methods in the Humanities and Social Sciences. Springer International Publishing, Berlin, 2018.

[HFM15] HINRICHS U., FORLINI S., MOYNIHAN B.: Speculative practices: Utilizing infovis to explore untapped literary collections. *IEEE Transactions on Visualization and Computer Graphics 22*, 1 (2015), 429–438. doi:10.1109/TVCG.2015.2467452.

[HGP99] HOFFMAN P., GRINSTEIN G., PINKNEY D.: Dimensional anchors: a graphic primitive for multidimensional multivariate information visualizations. In *Proceedings ACM Workshop on New Paradigms in Information Visualization and Manipulation* (1999), pp. 9–16. doi:10.1145/331770.331775.

[HLT*20] HUANG G., LI Y., TAN X., TAN Y., LU X.: Planet: A radial layout algorithm for network visualization. *Physica A: Statistical Mechanics and its Applications 539* (2020), 122948. doi:10.1016/j.physa.2019.122948.

[Hol06] HOLTEN D.: Hierarchical edge bundles: Visualization of adjacency relations in hierarchical data. *IEEE Transactions on Visualization and Computer Graphics 12*, 5 (2006), 741–748. doi:10.1109/TVCG.2006.147.

[HSS*20] HINTERREITER A., STEINPARZ C., SCHÖFL M., STITZ H., STREIT M.: Exploring visual patterns in projected human and machine decision-making paths. *arXiv preprint arXiv:2001.08372* (2020). doi:arXiv:2001.08372.

[HSSS17] HOHMAN F., SONI S., STEWART I., STASKO J.: A viz of ice and fire: Exploring entertainment video using color and dialogue. In *IEEE VIS Workshop on Visualization for the Digital Humanities* (2017), vol. 2.

[HVP*19] HÖLLT T., VILANOVA A., PEZZOTTI N., LELIEVELDT B. P. F., HAUSER H.: Focus+ context exploration of hierarchical embeddings. In *Computer Graphics Forum* (2019), vol. 38, pp. 569–579. doi:10.1111/cgf.13711.

[IM15] INGRAM S., MUNZNER T.: Dimensionality reduction for documents with nearest neighbor queries. *Neurocomputing 150* (2015), 557–569. doi:10.1016/j.neucom.2014.07.073.

[KK96] KEIM D. A., KRIEGEL H.-P.: Visualization techniques for mining large databases: A comparison. *IEEE Transactions on knowledge and data engineering 8*, 6 (1996), 923–938. doi:10.1109/69.553159.

[KRM*17] KRUIGER J. F., RAUBER P. E., MARTINS R. M., KERREN A., KOBOUROV S., TELEA A. C.: Graph layouts by t-sne. In *Computer Graphics Forum* (2017), vol. 36, Wiley Online Library, pp. 283–294.

[Lan95] LANT A.: Haptical Cinema. *October 74, Autumn 1995* (1995), 45–73.

[LLZ15] LIU J., LUGHOFER E., ZENG X.: Aesthetic perception of visual textures: a holistic exploration using texture analysis, psychological experiment, and perception modeling. *Frontiers in Computational Neuroscience 9* (2015), 134.

[LZD*19] LI G., ZHANG Y., DONG Y., LIANG J., ZHANG J., WANG J., MCGUFFIN M. J., YUAN X.: Barcodetree: Scalable comparison of multiple hierarchies. *IEEE Transactions on Visualization and Computer Graphics 26*, 1 (2019), 1022–1032. doi:10.1109/TVCG.2019.2934535.

[Mor13] MORETTI F.: *Distant Reading*. Verso Books, 2013.

[MvdGvdM*15] MAHFOUZ A., VAN DE GIESSEN M., VAN DER MAATEN L., HUISMAN S., REINDERS M., HAWRYLYCZ M. J., LELIEVELDT B. P. F.: Visualizing the spatial gene expression organization in the brain through non-linear similarity embeddings. *Methods 73* (2015), 79–89. doi:10.1016/j.ymeth.2014.10.004.

[NL12] NADOOLMAN LANDIS D.: *Hollywood Costume*. Harry N. Abrams Inc, New York, 2012.

[Ole17] OLESEN C. G.: *Film History in the Making. Film Historiography, Digitised Archives and Digital Research Dispositifs*. PhD Thesis, University of Amsterdam, Amsterdam, 2017.

[PHL*16] PEZZOTTI N., HÖLLT T., LELIEVELDT B., EISEMANN E., VILANOVA A.: Hierarchical stochastic neighbor embedding. In *Computer Graphics Forum* (2016), vol. 35, Wiley Online Library, pp. 21–30.

[PHS06] PAVLO A., HOMAN C., SCHULL J.: A parent-centered radial layout algorithm for interactive graph visualization and animation. arXiv:preprintcs/0606007, doi:arXiv:cs/0606007.

[Pla13] PLATZER A.: Visualization of SNPs with t-sne. *PloS one 8*, 2 (2013), e56883. doi:10.1371/journal.pone.0056883.

[PT19] PAGLIOSA L. D. C., TELEA A. C.: Radviz++: Improvements on radial-based visualizations. In *Informatics* (2019), vol. 6, Multidisciplinary Digital Publishing Institute, p. 16. doi:10.3390/informatics6020016.

[RFT16] RAUBER P. E., FALCÃO A. X., TELEA A. C.: Visualizing time-dependent data using dynamic t-sne. In *EuroVis (Short Papers)* (2016), pp. 73–77.

[Rob05] ROBERTS J. C.: Exploratory visualization with multiple linked views. In *Exploring geovisualization*. Elsevier, 2005, pp. 159–180.

[RPH20] ROBINSON I., PIERCE-HOFFMAN E.: Tree-sne: Hierarchical clustering and visualization using t-sne. *arXiv preprint arXiv:2002.05687* (2020). doi:arXiv:2002.05687.

[SGL08] STASKO J., GÖRG C., LIU Z.: Jigsaw: Supporting investigative analysis through interactive visualization. *Information visualization 7*, 2 (2008), 118–132. doi:10.1109/VAST.2007.4389006.

[SHS10] SCHULZ H.-J., HADLAK S., SCHUMANN H.: The design space of implicit hierarchy visualization: A survey. *IEEE Transactions on Visualization and Computer Graphics 17*, 4 (2010), 393–411.

[Stu16] STUTZ O. K.: *Algorithmische Farbfilmästhetik. Historische sowie experimentell-digitale Notations- und Visualisierungssysteme des Farbfilms im Zeichen der Digital Humanities 2.0 und 3.0*. Master's thesis, Universität Zürich, Zürich, 2016.

[SVI*16] SZEGEDY C., VANHOUCKE V., IOFFE S., SHLENS J., WOJNA Z.: Rethinking the inception architecture for computer vision. In *Proceedings IEEE Computer Vision and Pattern Recognition* (2016), pp. 2818–2826.

[VBW15] VEHLOW C., BECK F., WEISKOPF D.: The State of the Art in Visualizing Group Structures in Graphs. Borgo R., Ganovelli F., Viola I., (Eds.). doi:10.2312/eurovisstar.20151110.

[VDM14] VAN DER MAATEN L.: Accelerating t-sne using tree-based algorithms. *The Journal of Machine Learning Research 15*, 1 (2014), 3221–3245.

[VLKS*11] VON LANDESBERGER T., KUIJPER A., SCHRECK T., KOHLHAMMER J., VAN WIJK J. J., FEKETE J.-D., FELLNER D. W.: Visual analysis of large graphs: State-of-the-art and future research challenges. In *Computer graphics forum* (2011), vol. 30, pp. 1719–1749.

[vUHP*17] VAN UNEN V., HÖLLT T., PEZZOTTI N., LI N., REINDERS M. J. T., EISEMANN E., KONING F., VILANOVA A., LELIEVELDT B. P. F.: Visual analysis of mass cytometry data by hierarchical stochastic neighbour embedding reveals rare cell types. *Nature Communications 8*, 1 (2017), 1–10. doi:10.1038/s41467-017-01689-9.

[VWF09] VIEGAS F. B., WATTENBERG M., FEINBERG J.: Participatory visualization with wordle. *IEEE Transactions on Visualization and Computer Graphics 15*, 6 (2009), 1137–1144. doi:10.1109/TVCG.2009.171.

[ZJHJ16] ZUO H., JONES M., HOPE T., JONES R.: Sensory Perception of Material Texture in Consumer Products. *The Design Journal 19.03* (2016), 405–427.