



Collection Space Navigator: An Interactive Visualization Interface for Multidimensional Datasets

Tillmann Ohm *

mail@tillmannohm.com

School of Digital Technologies, Tallinn University
Tallinn, Estonia

Andres Karjus

School of Humanities, Tallinn University
Tallinn, Estonia

Mar Canet Sola *

mar.canet@tlu.ee

Baltic Film, Media and Arts School, Tallinn University
Tallinn, Estonia

Maximilian Schich

Baltic Film, Media and Arts School, Tallinn University
Tallinn, Estonia

ABSTRACT

We introduce the Collection Space Navigator (CSN), a browser-based visualization tool to explore, research, and curate large collections of visual digital artifacts that are associated with multidimensional data, such as vector embeddings or tables of metadata. Media objects such as images are often encoded as numerical vectors, based on metadata or using machine learning embeddings. Yet it remains a challenge to explore, analyze, and understand the resulting multidimensional spaces. Dimensionality reduction techniques such as t-SNE or UMAP often serve to project high-dimensional data into low dimensional visualizations, but require interpretation themselves given their typically abstract dimensions. The Collection Space Navigator provides a customizable interface that combines two-dimensional projections with an array of configurable multi-functional filters and navigation controls. The user is able to view and investigate collections by zooming and scaling, transforming between projections, and filtering dimensions via range sliders and text filters. Insights gained through these interactions can be used to augment original data via easy to use export capabilities. This paper comes with a functional online demo showcasing a large digitized collection of classical Western art. Users can reconfigure the interface to fit their own data and research needs, including projections and filter controls. This open source tool is intended to be applicable in a broad range of use cases, types of collections and across diverse disciplines.

CCS CONCEPTS

• **Human-centered computing** → **Visual analytics; Information visualization; Visualization toolkits.**

KEYWORDS

visualization, vector embeddings, metadata, interface, cultural data, datasets, text-to-image, information visualization

*Both authors contributed equally

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

VINCI 2023, September 22–24, 2023, Guangzhou, China

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0751-3/23/09.

<https://doi.org/10.1145/3615522.3615546>

ACM Reference Format:

Tillmann Ohm, Mar Canet Sola, Andres Karjus, and Maximilian Schich. 2023. Collection Space Navigator: An Interactive Visualization Interface for Multidimensional Datasets. In *The 16th International Symposium on Visual Information Communication and Interaction (VINCI 2023)*, September 22–24, 2023, Guangzhou, China. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3615522.3615546>

1 INTRODUCTION

Large collections of digital artifacts and associated metadata can be effectively studied using browsable interactive visualizations that reflect and resonate with the intrinsic shape of their data [20]. Mapping the topology of a collection into a multidimensional space can help to better understand the overall structure of a dataset and can uncover patterns hinting at underlying trends and dynamics. For example, a researcher or curator may visually explore the space, as constituted by some measure of artifact similarity, looking at different groups of similar objects to identify regions of interest for further quantitative and qualitative investigation.

Multidimensional feature vectors can be used to further describe artifact properties. This can include both categorical and numerical information. Numerical properties can be derived directly from metadata, such as in the case of artifact creation dates, or constructed through various feature extraction techniques. Neural network methods, for example, can encode measures of complex text semantics [8], of visual image properties [12, 15, 16, 33], of joint image-text pair embeddings [27, 29, 35], or of spectral audio features [28]. Hand-crafted feature engineering approaches [41] and algorithmic approaches such as compression ensembles [14] offer interpretable vector representations. Visualizing and exploring patterns in metadata such as artwork creation or acquisition dates can also be informative on its own [36].

Dimensionality reduction techniques can be used to reduce high-dimensional data to a more manageable number of dimensions by remapping or projecting the multidimensional topology into a lower-dimensional coordinate space [2, 22, 24, 37, 38]. For visual interpretation of multidimensional embedding spaces, such projection methods are used to present the data in two or three dimensions, which essentially can function as a reference topography of the original high-dimensional topology. The challenge for dimensionality reduction techniques is to preserve complex relationships, while necessarily compressing information: objects close to each other in the original space should ideally also be close in the low-dimensional topographic projection space.

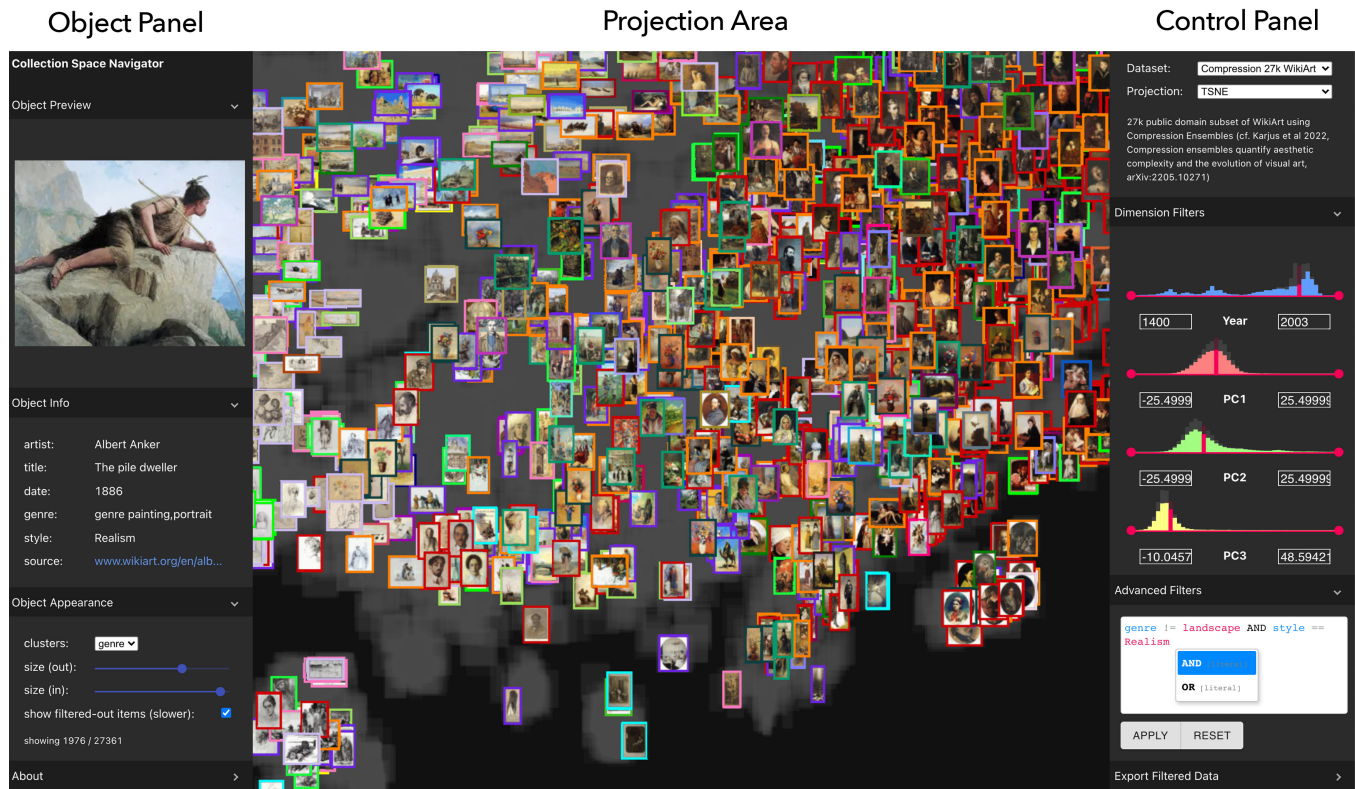


Figure 1: The Collection Space Navigator (CSN). The central *Projection Area* displays a x-y scatter plot of images based on the selected projection (e.g. UMAP, t-SNE), with filtered images greyed out, and mouse-over highlight. The *Object Panel* (left) shows a larger *Object Preview* of the highlighted image, together with *Object Info* based on selected metadata; *Object Appearance* visualizes clusters (optional), sets the projection thumbnail size (zoomed-out and zoomed-in). The *Control Panel* (right) allows for selection of *Data and Projections*; custom interactive *Dimension Filters* and *Advanced Filters* facilitate dataset exploration, analysis, and understanding (see text); the filtered object metadata and current projection view can be downloaded via *Export Filtered Data*.

Multiple projection views can help to better understand multidimensional data [2]. Two-dimensional static projections are visually comprehensible, but can only provide a limited view into multidimensional data. To gain intuition over a high-dimensional vector space, it can be helpful to interpret and compare many different projections. Interactive components in graphical user interfaces (GUI) can further help to gain intuition over the complex interactions between many dimensions. GUI elements such as range sliders are particularly useful to navigate through multiple features and dimensions, and to query and filter the data [1, 40].

2 RELATED WORK

Several interactive visualization experiments and prototypes have emerged in recent years which mediate high-dimensional embeddings through low-dimensional projections. Their focus vary: some aim to provide an intuitive understanding of multidimensionality and embedding methods by visualizing datasets commonly used for Machine Learning tasks [7, 13, 17, 34], or offer explorative interfaces

to similarity spaces of cultural collections [9–11, 42]. These interactive visualization projects aim to provide overview and deeper insight into their collections, tailored to specific datasets. The VIKUS Viewer [25, 26] offers a more general framework for exploring cultural collections. It allows not only to view a collection as a similarity map of its image embeddings, but also to dynamically filter metadata such as time and categories. The Selfexploratory of the Selficity project [21] similarly uses a number of range sliders to filter a multidimensional image dataset.

The CSN user interface aligns with classic conventions of cultural cartography and scholarly figure design, using a cartesian projection with an auxiliary index and call-out details [23]. It combines this with the modern paradigm of interactive figure design [5, 39], allowing for a deeper functional user experience (UX) and understanding of multidimensional data. The navigation paradigm of the CSN range sliders, which function as Dimension Filters, further resonates with the recent state-of-the-art of understanding mathematical multidimensionality via interactive animation [30]. The CSN combines these foundations with the paradigm of a scatter plot of images [19]. The CSN is also functionally similar to network

visualization applications, such as Cytoscape [31], Helios Web [32] or Gephi [3], which focus on depicting another (yet related) form of multidimensionality in node-link diagrams of complex networks [6].

Relatedly, the authors of the TensorFlow Embedding Projector [34] suggest to include multipanel projections, i.e. more than one simultaneous projection panel. This would also make sense as a possible extension of the CSN and would be in line with the prevalence of “multi-chart” figure panels in multidisciplinary science journals [4, 18].

3 THE COLLECTION SPACE NAVIGATOR

3.1 Motivation

We developed the CSN as a flexible browser-based research tool applicable across various use cases and research domains:

- (1) Researching large collections of digital objects (e.g. images, videos, audio, text, 3D models) with the ability to identify patterns and similar groups based on metadata and vector embeddings;
- (2) Understanding multidimensionality and projection methods by comparing different embedding spaces and dimensionality reduction techniques through intuitive navigation;
- (3) Presenting entire media collections online and communicating research findings with diverse audiences.

While prototypes and use cases exist for each of these aspects as discussed above, we are not aware of a tool that meets all of these requirements. Our contribution therefore lies in the combination and extension of existing interfaces to work towards a more universal, cross-platform, open modular research and curation system. It is highly optimized and capable of displaying large collections of hundreds of thousands of artifacts even on consumer hardware.

3.2 Design principles

To make complex interactions in multifaceted datasets comprehensible while meeting the diverse needs of users such as researchers and curators, we formulated three design principles for the tool:

- (1) Providing an open modular system that adapts to different research needs, domains and datasets while preventing information overload;
- (2) Providing a complete overview of the collection while encouraging immersive exploration of the objects;
- (3) Providing a multitude of interaction mechanisms and modalities to foster intuition, such as zooming, panning, hovering, and sliding through feature dimensions.

3.3 Components

3.3.1 Projection Area. The central part of the interface is the *Projection Area* (Figure 1 center). It maps the given input collection in its entirety as miniature images in an interactive 2D scatter plot, with coordinates defined by the chosen projection method (Figure 2). Basic navigation operations such as zooming or “drag and move” allow free exploration of the projection space. While the user sees only 2 dimensions in the central projection area, the CSN technically includes a third axis for depth. Moving along the depth axis (by




Figure 2: Examples of various 2D projections and visualization features in the CSN tool. a) UMAP projection with large thumbnails, providing a comprehensive view of the image content; b) UMAP projection with medium size thumbnails and cluster colors of categorical data selected in *Object Appearance*; c) UMAP projection with medium size thumbnails and filtered-out objects in grey; d) UMAP projection with small thumbnails and cluster colors, providing a more compact representation; e) t-SNE projection, showing an alternative dimensionality reduction technique; f) Simple x/y plot, here showing a principal component over time for temporal analysis. This flexible tool enables effective explorations and comparison of data across different visualization methods via multiple different 2D projections. These are flexible and based on data import configuration, selectable in *Data & Projections*. Thumbnail size and cluster highlights can be simultaneously adjusted via the *Object View Settings*.

zooming) effectively reveals overlapping objects. The appearance of the thumbnails can be adjusted in the *Object Panel*.

3.3.2 Object Panel. The *Object Panel* (Figure 1 left) has three collapsible sub-menus: *Object Preview*, *Object Info* and *Object Appearance*. The *Object Preview* section displays a larger version of the miniature thumbnail currently hovered on in the *Projection Area*. By default it simply shows a larger version of the same thumbnail, but it can be set to display a higher resolution version of the hovered image, stored either locally or remotely.

The *Object Info* section provides detailed information on the currently selected object. This aspect of the CSN is highly flexible, as the metadata fields that provide this information can be easily defined in the separate configuration file. Minimally it can display just the file name, but it can equally well include extensive metadata – for example in the case of art collections, the author, production year, location, genre, style, and other details.

The *Object Appearance* section contains options to control the visual appearance of the objects in the *Projection Area*. A predefined group (from categorical metadata) can be selected from a drop-down list to show clusters. Objects of the same category are depicted with the same color border around their thumbnails. The size and scale of the thumbnail images are adjustable with convenient sliders. Size

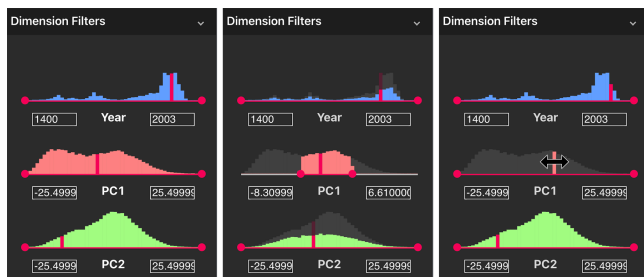


Figure 3: Interactive Dimension Filters. Left: Unfiltered Dimension Filters, consisting of range sliders with interactive histograms above them, showing the distribution of all objects along the slider’s dimension with the bin of the currently selected object highlighted in red. Center: reducing the range of one slider affects the distribution of all dimensions, reflected by the histograms. Right: *Bin Mode* functionality is activated by clicking on a histogram, allowing the user to activate one bin at a time, with the Projection Area displaying the corresponding objects within the active bin. These interactive Dimension Filter features allow in-depth exploration and visualization of multi-dimensional data distributions to gain a deeper understanding of the relationships between data points.

determines how large the thumbnails should be when fully zoomed out, while Scale affects the size when zoomed in.

3.3.3 Control Panel. The *Control Panel* (Figure 1 right) has four collapsible sub-menus: *Data & Projections*, *Dimension Filters*, *Advanced Filters*, and *Export Filtered Data*. The *Data & Projections* section contains a *Dataset* drop-down list of selectable datasets. For very large collections, we recommend providing a smaller subset by default and offering the entire set on demand (such subsets can be conveniently produced using the CSN configuration Python notebook). The section also contains a *Projection* drop-down list of selectable projections and mappings. Switching between different projections, e.g. different embedding or reduction methods, smoothly rearranges the positions of the objects in the *Projection Area*. These intuitive animations can provide new insights into the intermediate state between two projections and expose their differences.

The *Dimension Filters* are optional interactive elements that function to filter the objects in the *Projection Panel* (Figure 3). They control the range of the assigned variables, which could be dimensions of the embedding, metadata such as date or year of creation of an artwork, or inferred properties of the image such as colorfulness or contrast. Histograms above the range sliders provide additional statistical information and feedback on how changes affect the distribution of the mapping. They are constantly updated to reflect the distribution of the entire dataset as well as the distribution of filtered and unfiltered objects. Clicking on a histogram activates the *Bin Mode*: moving the cursor over the bars of the histogram temporarily displays only objects within the narrow range of the bar. A second click terminates this function and sets the filters back to the previous state. Additionally, hovering over the thumbnails

in the *Projection Area* highlights the corresponding vertical bar in each histogram.

The *Advanced Filters* section is a text field to construct and apply search and filter queries. By default, it handles basic query operators such as AND, OR, equals ($=$), does not equal (\neq), as well as custom operators. Nested and complex search queries are also supported using round brackets. When setting up the CSN for a new collection, using the configuration file (in Jupyter notebook format), each metadata field can be defined as a Free Text Entry (enabling queries) or as a Categorical Selection (generating drop-down lists, i.e. GUI elements that allow simple search and selection). The *Export Filtered Data* section in the *Control Panel* allows downloading the metadata of the currently filtered objects as a CSV file, and the current projection view as PNG file.

4 CONCLUSION

The CSN is a flexible and powerful tool for visualizing large multi-dimensional datasets, including embedding vectors and metadata. Available publicly on Github, the CSN includes a demo and a usage guide via an interactive Jupyter notebook. It’s adaptable to a range of research and data needs, supporting various vector types, metadata, projection methods, and representations, in a user-friendly interface. Though we demonstrate CSN using visual data, it’s not limited to image-derived vectors. It can equally visualize embeddings or metadata from audio or text, represented by appropriate thumbnails or labels, to explore non-visual meaning spaces. The CSN is open-source, and we encourage development and extensions. In conclusion, the CSN is a research tool for exploring, studying and curating large digital artefact collections that supports multidisciplinary research and understanding of multidimensional meaning spaces.

CODE AVAILABILITY AND DEMO

The CSN is released as MIT license with code and documentation available at <https://github.com/Collection-Space-Navigator/CSN>. A live demo, using the same example data as in the paper figures, is available at <https://collection-space-navigator.github.io/CSN>

ACKNOWLEDGMENTS

T.O. and M.C. designed, co-authored, and developed the Collection Space Navigator software. T.O., M.C., A.K. and M.S. contributed to the research design and co-wrote the manuscript. T.O., M.C. and A.K. collected data. T.O. and M.C. contributed equally to this work as first authors. The authors thank Sebastian Ahnert, Mila Oiva, and the entire CUDAN team for useful conversations and input. All authors are supported by the CUDAN ERA Chair project, funded through the European Union’s Horizon 2020 research and innovation program (Grant No. 810961).

REFERENCES

- [1] Christopher Ahlberg and Ben Shneiderman. 1994. Visual information seeking using the filmfinder. In *Conference companion on Human factors in computing systems*. 433–434.
- [2] Ehsan Amid and Manfred K Warmuth. 2019. TriMap: Large-scale dimensionality reduction using triplets. *arXiv preprint arXiv:1910.00204* (2019).
- [3] Mathieu Bastian, Sebastien Heymann, and Mathieu Jacomy. 2009. Gephi: an open source software for exploring and manipulating networks. In *Proceedings of the international AAAI conference on web and social media*, Vol. 3. 361–362.

- [4] Heinrich Karl Wilhelm Berghaus. 1845. *Physikalischer Atlas oder Sammlung von Karten: auf denen die hauptsächlichsten Erscheinungen der anorganischen und organischen Natur nach ihrer geographischen Verbreitung und Vertheilung bildlich dargestellt sind*. Text. Vol. 1. Justus Perthes.
- [5] Jacques Bertin. 2010. Semiology of graphics: Diagrams, networks, maps.
- [6] Jan Niklas Böhm, Philipp Berens, and Dmitry Kobak. 2022. Attraction-repulsion spectrum in neighbor embeddings. *Journal of Machine Learning Research* 23, 95 (2022), 1–32.
- [7] Grant Custer. 2020. An interactive UMAP visualization of the MNIST data set. <https://github.com/GrantCuster/umap-explorer>. [Online; accessed 10-May-2023].
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [9] Cyril Diagne, Nicolas Barradeau, and Simon Doury. 2018. t-SNE Map. <https://experiments.withgoogle.com/t-sne-map>. [Online; accessed 19-July-2022].
- [10] Douglas Duhaime. 2017. PixPlot. <https://github.com/YaleDHLab/pix-plot>. [Online; accessed 19-July-2022].
- [11] Katrin Glinka, Christopher Pietsch, and Marian Dörk. 2017. Past Visions and Reconciling Views: Visualizing Time, Texture and Themes in Cultural Collections. *DHJ: Digital Humanities Quarterly* 2 (2017).
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [13] Zeyang Huang, Daniel Witschard, Kostiantyn Kucher, and Andreas Kerren. 2023. VA+ Embeddings STAR: A State-of-the-Art Report on the Use of Embeddings in Visual Analytics. In *COMPUTER GRAPHICS forum*, Vol. 42.
- [14] Andres Karjus, Mar Canet Solà, Tillmann Ohm, Sebastian E Ahnert, and Maximilian Schich. 2023. Compression ensembles quantify aesthetic complexity and the evolution of visual art. *EPJ Data Science* 12, 1 (2023), 21.
- [15] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. 2020. Big transfer (bit): General visual representation learning. In *European conference on computer vision*. Springer, 491–507.
- [16] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2017. Imagenet classification with deep convolutional neural networks. *Commun. ACM* 60, 6 (2017), 84–90.
- [17] Cloudera Fast Forward Labs. 2019. Active Learner prototype. <https://activelearner.fastforwardlabs.com>.
- [18] Po-shen Lee, Jevin D West, and Bill Howe. 2017. Viziometrics: Analyzing visual information in the scientific literature. *IEEE Transactions on Big Data* 4, 1 (2017), 117–129.
- [19] Lev Manovich. 2012. How to compare one million images? In *Understanding digital humanities*. Springer, 249–278.
- [20] Lev Manovich. 2012. Media visualization: Visual techniques for exploring large media collections. *Media studies futures* (2012), 1–21.
- [21] Lev Manovich, Moritz Stefaner, Mehrdad Yazdani, Dominikus Baur, Daniel Goddmeier, Alise Tifentale, Nadav Hochman, and Jay Chow. 2015. The Selfexploratory. <https://selfiecity.net/selfieexploratory/>. [Online; accessed 19-Dec-2022].
- [22] Leland McInnes, John Healy, and James Melville. 2018. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426* (2018).
- [23] Giambattista Nolli. 1748. *La nuova topografia di Roma*. Rome.
- [24] Karl Pearson. 1901. LIII. On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin philosophical magazine and journal of science* 2, 11 (1901), 559–572.
- [25] Christopher Pietsch. 2017. VIKUS Viewer. <https://github.com/cpietsch/vikus-viewer>. [Online; accessed 19-July-2022].
- [26] Christopher Pietsch, Viktoria Brüggemann, Mark-Jan Bludau, and Marian Dörk. 2020. Visuelle Exploration Zweier Musealer Sammlungen. <https://visualisierung.smb.museum>. [Online; accessed 19-July-2022].
- [27] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*. PMLR, 8748–8763.
- [28] Zhao Ren, Kun Qian, Zixing Zhang, Vedhas Pandit, Alice Baird, and Bjorn Schuller. 2018. Deep scalogram representations for acoustic scene classification. *IEEE/CAA Journal of Automatica Sinica* 5, 3 (2018), 662–669.
- [29] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10684–10695.
- [30] Grant Sanderson. 2017. Thinking outside the 10-dimensional box. <https://www.youtube.com/watch?v=zwAD6dRSVyI>. [Online; accessed 19-December-20022].
- [31] Paul Shannon, Andrew Markiel, Owen Ozier, Nitin S Baliga, Jonathan T Wang, Daniel Ramage, Nada Amin, Benno Schwikowski, and Trey Ideker. 2003. Cytoscape: a software environment for integrated models of biomolecular interaction networks. *Genome research* 13, 11 (2003), 2498–2504.
- [32] Filipi Nascimento Silva. 2023. Helios Web. <https://github.com/filipinascimento/helios-web>. [Online; accessed 10-July-2023].
- [33] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
- [34] Daniel Smilkov, Nikhil Thorat, Charles Nicholson, Emily Reif, Fernanda B Viégas, and Martin Wattenberg. 2016. Embedding projector: Interactive visualization and interpretation of embeddings. *arXiv preprint arXiv:1611.05469* (2016).
- [35] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*. PMLR, 2256–2265.
- [36] Mar Canet Solà, Antonina Korepanova, Ksenia Mukhina, and Maximilian Schich. 2023. Quantifying collection lag in European modern and contemporary art museums. *arXiv preprint arXiv:2305.14159* (2023).
- [37] Jian Tang, Jingzhou Liu, Ming Zhang, and Qiaozhu Mei. 2016. Visualizing large-scale and high-dimensional data. In *Proceedings of the 25th international conference on world wide web*. 287–297.
- [38] Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, 11 (2008).
- [39] Bret Victor. 2011. Scientific communication as sequential art. <http://worrydream.com/ScientificCommunicationAsSequentialArt>. [Online; accessed 10-May-2023].
- [40] Christopher Williamson and Ben Shneiderman. 1992. The Dynamic HomeFinder: Evaluating dynamic queries in a real-estate information exploration system. In *Proceedings of the 15th annual international ACM SIGIR conference on Research and development in information retrieval*. 338–346.
- [41] Jiajing Zhang, Yongwei Miao, Junsong Zhang, and Jinhui Yu. 2020. Inkthetics: a comprehensive computational model for aesthetic evaluation of Chinese ink paintings. *IEEE Access* 8 (2020), 225857–225871.
- [42] Audun M. Øygard. 2017. Visualizing an Art Collection. <http://auduno.github.io/2018/10/27/visualizing-an-art-collection/index.html>. [Online; accessed 19-July-2022].