Theme Article: Art and Cultural Heritage

Many Views Are Not Enough: Designing for Synoptic Insights in Cultural Collections

Florian Windhager
Danube University Krems

Saminu Salisu
Danube University Krems

Roger A. Leite
Vienna University of Technology

Velitchko Filipov
Vienna University of Technology

Silvia Miksch
Vienna University of Technology

Günther Schreder
Danube University Krems

Eva Mayr
Danube University Krems

***Abstract*—Cultural object collections attract and delight spectators since ancient times. Yet, they also easily overwhelm visitors due to their perceptual richness and associated information. Similarly, digitized collections appear as complex, multifaceted phenomena, which can be challenging to grasp and navigate. Though visualizations can create various types of collection overviews for that matter, they do not easily assemble into a "big picture" or lead to an integrated understanding. We introduce coherence techniques to maximize connections between multiple views and apply them to the prototype PolyCube system of collection visualization: with map, set, and network visualizations it makes spatial, categorical, and relational collection aspects visible. For the essential temporal dimension, it offers four different views: superimposition, animation, juxtaposition, and space–time cube representations. A user study confirmed that better integrated visualizations support synoptic, cross-dimensional insights. An outlook is dedicated to the system's applicability within other arts and humanities data domains.**

Digital Object Identifier 10.1109/MCG.2020.2985368
Date of publication 6 April 2020; date of current version 28 April 2020.

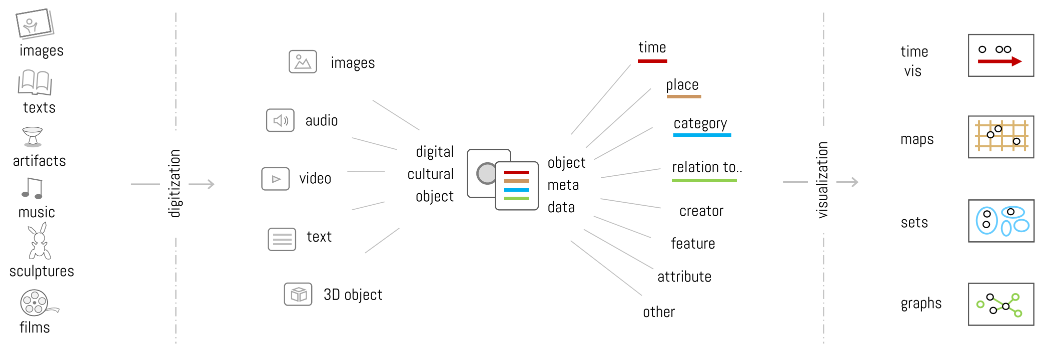

Figure 1. Digital cultural objects (center) are frequently connected to multiple dimensions of metadata, which can be encoded into different aggregated views of a collection (right).

Cultures collect objects, achievements, and practices deemed worthy of preservation due to their aesthetic, historic, scientific, or social value. Modern societies thus incorporate a whole system of GLAM institutions (i.e., galleries, libraries, archives, and museums) dedicated to these preservation practices and to the purpose of collective learning and intergenerational transmission. Recently, digitization is providing new ways and means for this endeavor, but the basic cognitive challenge for visitors and spectators stays the same: Cultural collections are immensely rich object assemblies of high perceptual complexity, which are further connected to large amounts of additional information and which are frequently condensed into relatively small exhibition spaces. Computer screens—whether stationary or mobile—make no exception. Visitors to such information spaces thus are easily overwhelmed by an abundance of details, and can swiftly feel "museum fatigue" due to the huge amount of information,[1] which is why various techniques have been developed to provide them with a collection overview in advance, which assists orientation and exploration.

In this context, visualization plays an essential role for *digital* cultural collections. In addition to realistic object images, which provide detailed views for selected objects (such as photographs or 3-D scans), information visualization provides manifold options to create collection overviews and browsable arrangements on screens.[2] Among other techniques, maps, set, and network visualizations play a frequent role, to convey visitors with orientation about origins, attributes, and relations within large corpora of objects. Also historical information (e.g., presented via timelines) plays an essential role in cultural heritage visualization, and it has been argued that a temporal perspective (i.e., a "diachronic view") could be productively added to every "synchronic" (i.e., nontemporal) visualization technique.[3]

The standard technique to actually draw together a combination of multiple perspectives is the use of *coordinated multiple views*.[4] Multiple views can effectively connect different perspectives into an ensemble of complementary composites, but they also come with specific downsides: They split the users' attention between different views, they increase the perceptual complexity, and their heterogeneous mapping techniques require the coordination of heterogeneous sense-making procedures. Consequently, the information provided by multiple views cannot easily be cognitively integrated.[5] Against this background, we argue that many views are good to have, but not good enough, if users should gain a better integrated, bigger picture of a complex phenomenon (see Box 1). With this article, we introduce a prototype visualization system for cultural collections doing just that: It unfolds a variety of views, but also goes beyond this essential standard technique to visualize information in a more integrated manner, to facilitate synoptic insights, and to foster cross-dimensional reasoning. In parallel, the underlying PolyCube project seeks to develop synoptic visualizations, which help users to develop an integrated understanding of multiple dimensions in the collection (a *mental model*) rather than a poorly connected mashup of information (a cognitive collage).* In the following, we zoom in on

*Online: http://www.donau-uni.ac.at/en/polycube/

BOX1: SOME REMARKS ON COGNITION

The design rationale of the presented system draws on basic cognitive assumptions. For a better understanding, we summarize some of the guiding concepts deemed relevant for reasoning with multiple views. For a more detailed discussion, see Windhager *et al.*[6,7]

Mental models: In cognitive terms, while users explore and observe an information space, they continuously construct an internal representation of this space and use this so-called mental model to interpret the information and to explore it. Mental models are not fixed, but are dynamic representations that develop over time, growing richer in structure and content.

Coherence: When confronted with a task or problem, users utilize and inspect their mental model to find and evaluate possible solutions. In these cases, internal representations are more effective which connect their different pieces of information in a coherent and consistent fashion. In contrast to such a well-connected mental model, distorted and loosely connected internal representations were termed *cognitive collages*.

Integration: If complex conceptual entities and phenomena are visualized across multiple views, users will build up representations for each of these views, which are only loosely connected. To develop a coherent bigger picture—and to solve problems across multiple dia- and synchronic data dimensions—these separate information pieces have to be (re-)connected. This synthetical task turns out to require high amounts of cognitive load, as cognitive science research shows.[6] Therefore, to solve cross-dimensional problems with less effort, users require support to build up internal representations, which combine and interconnect multiple data dimensions consistently. This is the background against which the PolyCube project investigates design and coherence techniques—to go beyond multiple views and help users build up more *integrated* mental models.

challenges arising from the use of coordinated multiple views and delineate design rationales and specific "coherence techniques" for creating more integrated visualizations. Then, we introduce the PolyCube prototype, summarize the results of an evaluation, and sketch out future work.

BEYOND MULTIPLE VIEWS

The data structure of cultural objects can be complex (see Figure 1): In addition to realistic object images, collectors and curators commonly assemble a whole range of object metadata—and various datafication procedures can extract a whole range of additional object features.[2] Digital cultural collections thus commonly appear as complex, multidimensional datasets, where each object is connected to multiple metadata entries.

Given this informational richness, visualization designers commonly conclude that *"one view is not enough."*[3] Single visualization techniques can only make a certain selection of data dimensions visible. Multiple views, by contrast, multiply the analytical vistas or interpretive windows into the data and thus can help to maximize insight, balance the strengths and weaknesses of individual views, and avoid misinterpretation.[8] Most interfaces to cultural collections thus make use of this design strategy (e.g., by combining a geographic map with a faceted timeline), to make *more than one* data dimension visible at once and to offer more than just one dimension for conceptual orientation.[2]

Multiple views can also be a way to multiply insights for data dimensions with specific relevance. Time arguably is such an essential dimension in cultural heritage and history data and it has been said that it could be productively connected with every other "synchronic" visualization technique.[3] Where time plays a central role, systems thus can offer more than one encoding for the "diachronic" data dimension, so that strengths and weaknesses of specific time-visualizations can mutually complement and balance each other.[8]

Problems With Multiple Views

However, coordinated multiple views also come with notable downsides and cognitive

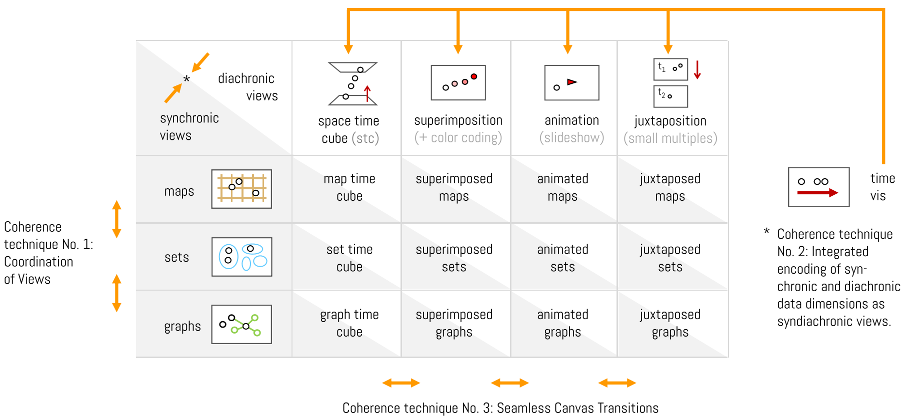

Figure 2. Cross-tabulation of the available synchronic views (lines) and diachronic views (columns) of the PolyCube system, together with annotation of implemented visual "coherence techniques" (orange).

overheads: They increase perceptual complexity and raise the efforts required to learn a system, while splitting user's attention during comparison of views and context switching.[9] Multiple views merge heterogeneous visualization techniques within one larger frame of display space, yet they cannot be directly integrated on a synoptic reasoning level (cf., the fable of the elephant and the blind men).[5] Juxtaposed information visualizations do not easily assemble into a "bigger picture" of the cultural collection, and do not necessarily lead to an integrated and synoptic understanding—except users invest high mental effort to reintegrate them on a nonperceptual level into one mental model (see Box 1[6,7]). Even then, there is no natural interconnectivity of different types of diagrams: They do not work like pieces of a puzzle or photographic perspectives, which can be easily reassembled to an integrated whole. If a "bigger picture" is intended, visualization designers have to actually support its creation—and interface designers have to invest additional efforts within the "gutter" of multiple views.[10]

Design Rationales

With specific regard to these challenges, we developed the PolyCube framework to investigate better integrated and more coherent visualization techniques. With focus on cultural collections, the project has been organized around a number of design requirements, which aim to foster synoptic insights and cross-dimensional reasoning. In short order, the PolyCube framework provides the following:

- multiple synchronic collection overviews,
- multiple types of diachronic overviews,
- access to single objects (close-up views), and
- flexible recombination of views, while
- maximizing visual system integration by the use of multiple *"coherence techniques."*

Figure 2 shows conceptually how the PolyCube framework interconnects the four collection visualization techniques, which have been outlined in Figure 1 (right-hand side).

Coherence Techniques

To develop a better integrated and coherent representation of the collection, we combine three distinct design strategies, which we refer to as *"coherence techniques,"*[7] i.e., techniques which help users to connect perspectives and information for the construction of a better integrated mental model. 1) We utilize the standard coherence technique of *coordination of views*,[4,9] which juxtaposes and interconnects a map, a set, and a network perspective via the interaction techniques of brushing, linking, and coordinated highlighting. 2) To go beyond multiple views, we resolve one standalone view—the one on time (cf. Baldonado *et al.*[9]: Rule of Parsimony)—and systematically integrate it into all other perspectives: While maps, sets, and networks commonly appear as *synchronic* (i.e., nontemporal or aggregated) views, we do not only enrich them with temporal information (i.e., coherence technique No. 2: "*integrated encoding*"), but offer four different ways to do so. As such, we "upgrade" each synchronic view to four types of *syn-diachronic* (i.e., time-oriented) views: i) a space–time cube representation (STC), ii) a color-coded superimposition view (SI), iii) an animated view (ANI), and iv) a juxtaposition view (JP). 3) Finally, to visually integrate these multiple diachronic views, we mediate the transitions between them with the coherence technique of "*animated canvas transitions*," which utilize dynamic space–time cube operations.[11] In joint, this package of coherence techniques aims to go beyond the common, casual bricolage of multiple views, and to provide a synoptic, but flexible system for complex collection visualization with a focus on the historical data dimension and visual and cognitive information integration.

Data, Users, Tasks

In order to further specify the design of our visualizations, we considered characteristics of data, users, and users' tasks in the cultural collection context.

Data: As we aimed to develop a visualization framework applicable to a wide range of cultural collections, we used two different datasets to inform and test our prototype design, which focus on photographs and films. The first dataset is a random selection of 2000 objects from the photography collection of Charles W. Cushman.[†] Cushman was a well-travelled American amateur photographer, who documented his activities by means of a travelogue, annotating all his photographs with locations, dates, keywords, and short descriptions. Each photographic object thus has categorical, temporal, and spatial metadata entries, and we additionally computed interobject relations based on keyword and description similarity.

The second dataset has been extracted from the Internet Movie Database (IMDB) by Spitz and Horvát[12] for over 55 000 movies, comprising information on genres, locations of production, year of release, and references to other movies (such as "is referenced," "is featured," "is version of," "is a sequel," "is a spin-off"). We reduced the dataset to keep only the most influential movies (i.e., which have been referenced at least 15 times in other movies), which reduced the dataset to approximately 2000 movies. Our IMDB dataset thus also features categorical, spatial, temporal, and relational metadata entries for each movie, together with a short description and a movie poster, which we added from the Open Movie Database (OMDb).[‡]

Users: Regarding the users of our system, we intended to support casual users, i.e., a heterogeneous group of users with different levels of visual literacy, expert knowledge, and interest. As they explore cultural collections mainly for leisure purposes, they are not necessarily motivated to invest high amounts of cognitive load to fully process and interconnect all information and perspectives available. Therefore, casual users can benefit from visualization designs, which actively assist them in the construction of a mental model.

In addition, it is important to be aware that casual users will only persist in exploration as long as it is rewarding for them (i.e., engaging, aesthetically pleasing, interesting, or intriguing). Hermetic and complex system designs can easily terminate the interaction at an early stage.

Tasks: Casual users of cultural collections often have no specific information needs, but look around and browse for something interesting, which they can access for details on demand. They do not pursue concrete tasks, but they are keen on gaining an overview and exploring the digital collection. When designing and developing PolyCube, we attempted to clarify these rather vague tasks as follows: 1) gaining a (synoptic) overview and conceptual orientation regarding the distribution of the major data dimensions of a cultural collection (e.g., time, space, categories, relations), 2) finding single objects of personal interest and inspecting their details, and 3) browsing through objects (e.g., according to time, relations, geographic origins, or shared categories). These tasks align with the task typology proposed by Brehmer and Munzner,[13] as the users in this case are consumers of the visualization and want to *explore*, *browse*, and *enjoy*.

[†] Online: https://webapp1.dlib.indiana.edu/cushman/, prepared by http://miriamposner.com/blog/getting-started-with-palladio/.

[‡] Online: http://www.omdbapi.com/

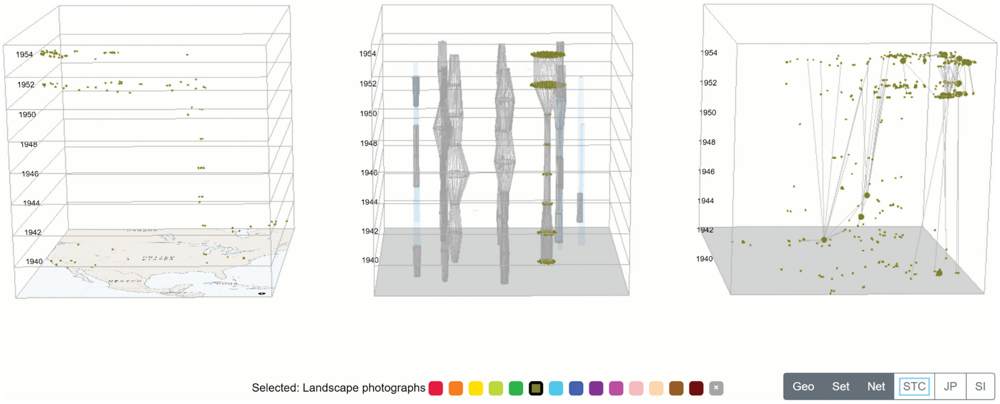

Figure 3. PolyCube system in its initial setup, providing a map, a set, and a network visualization of the Cushman photography collection from a space–time cube perspective, with one category filtered.

SYSTEM DESIGN

By its name, the PolyCube system refers to its initial syn-diachronic setup, which combines multiple space–time cube representations as “coordinated multiple cubes.” Figure 3 depicts this initial setup, which offers a multidimensional, time-oriented overview for different dimensions of a cultural collection. From left to right, it links a map-based view with a visualization of set-typed information and a graph-based representation of the collection (Coherence technique No. 1: coordination of views). Instead of following a common design practice and adding a coordinated timeline visualization as an extra view, STC representations encode diachronic and synchronic data dimensions in a visually and locally integrated fashion (coherence technique No. 2: integrated encoding).

Coordinated Multiple Cubes

A *map-based space–time cube* visualizes the spatio-temporal origins of a cultural object collection as a 3-D point cloud, where each object can be selected to access the detailed information and a preview. While the horizontal data plane shows a geographic map, the vertical axis represents time as an upward-pointing, spatial dimension. Temporal object positions can be visualized in a continuous or aggregated fashion (according to the chosen temporal granularity, signified by the number of horizontal slices). The resulting distribution of data points discloses the temporal and spatial origins and extension of a collection, together with spatio-temporal clusters and outliers.

A *set-based space–time cube* makes categorical or set-typed collection information visible by the means of circular set diagrams, in which data points, again, represent individual objects. Data points are horizontally arranged according to their primary set-affiliation (which can signify cultural styles, movements, curatorial categories, or genres), horizontally in time layers, and within each set-instance over time in a sunflower Phyllotaxis layout.[14] A visual hull structure can be activated to connect different temporal instances of each set, making set dynamics within a collection (like emergence, growth, diminution, breaks, or decline) visible and traceable over time.[5]

Finally, a *network-based space–time cube* represents relations between objects, based on a force-directed network layout. To avoid visual clutter generated by too many ties, we decided to cap the maximum number of displayed relations—or to hide links on demand. In case of weighted ties (Cushman), only the top similarity relations between objects are displayed as links. Also for nonweighted networks, only one random relation is visualized at the moment. Relations between objects thus form a 3-D meshwork, disclosing node clusters and outliers, as well as bundles and structures of relations. Structural node measures can be visualized on demand as glyph size (total degree centrality, as well as in-degree and out-degree), to easily identify the most important (e.g., most referenced) cultural objects.

To reduce the complexity for casual users and to accommodate also for simpler datasets, each space–time cube can be hidden on demand. Thereby, the PolyCube system offers multiple

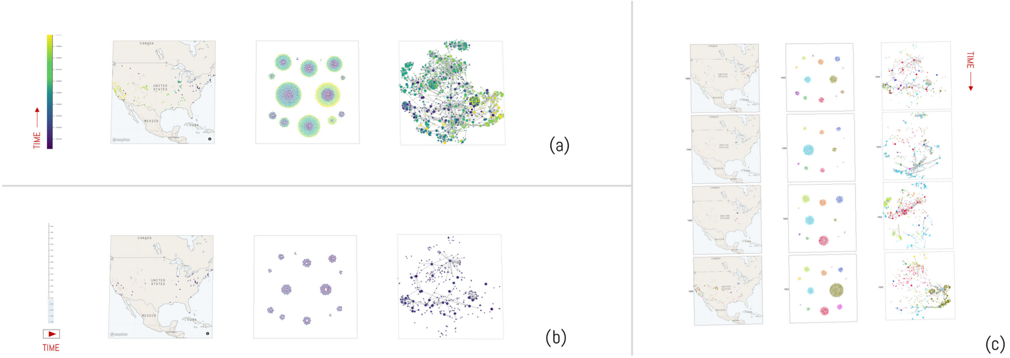

Figure 4. Three alternative syn-diachronic views support the visual analysis of the temporal developments of maps, sets, and networks by the means of (a) color coded superimposition, (b) animation, and (c) juxtaposition.

independent, but coordinated views that still sum up into a bigger picture of juxtaposed components when needed.

Further Syn-Diachronic Views

From a comparative point of view, all visualization techniques have different analytical strengths and weaknesses. This also applies for different visualizations of time[15,16]—and it holds true also for space–time cube encodings, which can suffer from visual clutter and other shortcomings (see more in the discussion). Due to the relevance of time for archival and historical collections, we thus provide multiple syn-diachronic views, to investigate the (geographical, set-typed, and relational) temporality of collections with a plurality of encoding techniques. To preserve the users' mental map during the switching of syn-diachronic views and to help them to integrate new perspectives into their prior constructed mental models, we utilize animated space–time cube operations as canvas transitions between the different syn-diachronic views (i.e., coherence technique No. 3).[11]

A *color-coded superimposition* view (SI) literally flattens the STC perspective, to offer an accumulated, top-down view on maps, sets, and networks. In this view, time is encoded by the use of a Viridis color scale, which also appears as an annotated legend on the left side of the screen (see Figure 4(a)).

An *animation function* (ANI) allows users to let a selected temporal brush traverse the dataset along the time axis (see Figure 4(b)). This animation function can be activated in the superimposition view, resulting in a flat animation, but it can also be activated in the two other syn-diachronic perspectives (STC and JP).

Finally, a *juxtaposed view* (JP) presents time in a "sliced" manner, arranging flattened time sections as panels side by side, from top to bottom (see Figure 4(c)). The number of panels is controlled by the users—and corresponds to the number of temporal slices shown in the STC representation.

Interaction Design

Many settings of the system can be modified and regulated by its users. Most fundamentally, users decide (via buttons at the bottom right) if they want to see all three synchronic views (maps, sets, networks)—or only a selection thereof. They also choose the syn-diachronic perspective of their liking (STC, SI, ANI, JP) and every chosen view can be zoomed and panned—and in case of the STC also rotated. When switching between syn-diachronic views (more specifically between STC, SI, and JP), the animated canvas transitions come into play. Figure 5 shows how these canvas transitions mediate all switching activities, so that the mental maps—which have already been created—can be preserved.

Object preview: Across all syn-diachronic views, clicking a data point opens a close-up-panel, presenting an object preview on the left side. The panel also displays selected object metadata (object categories, date and place of origin, degree centrality), similar objects, and a browsing function, which allows the user to swipe chronologically through a cultural collection.

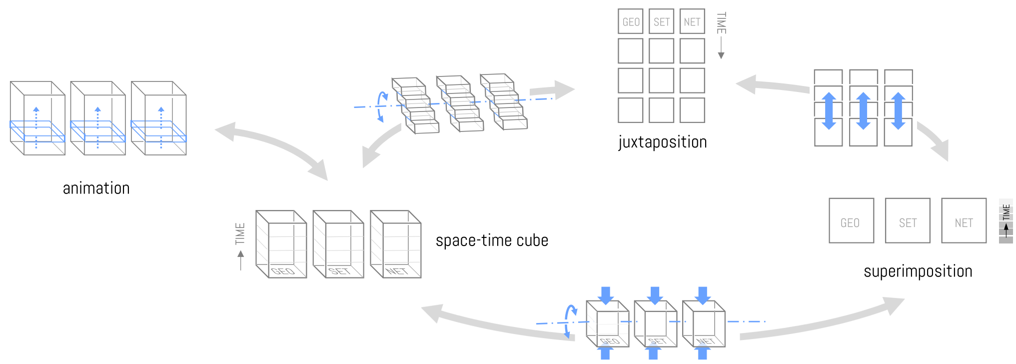

Figure 5. Animated canvas transitions between syn-diachronic views, based on space–time cube operations.

Temporal filtering: A timeline widget at the left side of the screen allows the user to filter and brush dynamically within the temporal scope of a collection. This functionality works across all syn-diachronic views in a coordinated fashion.

Glyph color: The system provides three color options for the data points: a uniform color (red), distinct categorical colors according to the set-assignment of objects (as indicated by a legend at the bottom of the screen), and a temporal color-coding, utilizing the Viridis color scale.

Noverlap function: With increasing size and spatio-temporal collection density, the placement of data points on maps can generate significant overlays, merging hundreds of objects into one glyph. As a clutter management technique, a user-controlled "noverlap"-function adds positional jitter to object positions in all map-based views, thus spreading out glyphs around their collective center point, which allows us to disentangle colocations and supports visual disambiguation.

Implementation Details

The PolyCube system is an open-source web application, and a demonstration version showcasing the data of the two case studies is available at the following URL: https://danubevislab.github.io/polycube/cga2020/. The source code and implementation along with steps on how to set up the system can be found at https://github.com/danubevislab/polycubeViews. In our implementation, we use D3.js (https://d3js.org/) to compute the layout of the set-based and network-based space–time cubes, three.js (https://threejs.org/) to create a 3-D environment and plot our data points, and angular (https://angular.io/) to manage the interactions, views, and navigation. For the map-based space–time cube, we also utilize Mapbox (https://www.mapbox.com/) to arrange our data points according to their geographic position.

EVALUATION

In the course of the system development, we conducted several user studies (for geotemporal views,[3,17] for set-temporal views[14]) and regular heuristic inspections. For the summative evaluation of the whole prototype system, we set up a qualitative study with ten casual users (under 30 yrs: 2; 30–40 yrs: 5; over 50 yrs: 3/university degree: 6/ female: 9). Their experience with visualization was mixed (5 low, 3 medium, 2 high), as was their interest in the domains of the collections—photography (0 low, 2 medium, 8 high) and films (1 low, 2 medium, 7 high). The test procedure took between 21 and 33 min ($M = 28$).

The aim of the evaluation was to understand how the coherence techniques could support casual users. More specifically, we focused on 1) how animated canvas transitions (coherence technique No. 3) supported the users' understanding of the different views, and 2) how the integrated encoding (coherence technique No. 2) and the coordination of views (coherence technique No. 1) affected users' synoptic insight into the collection. In addition, we inspected the overall usability of the system, and asked for user preferences regarding various syn-diachronic perspectives.

Procedure: Each participant tested the prototype individually on a 24" computer screen. In phase I, the experimenter started with a guided walk-through of the Cushman photography

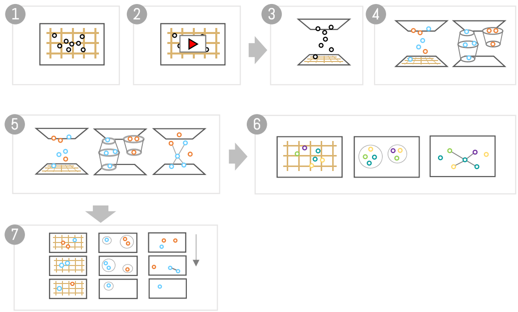

Figure 6. Guided walk through the PolyCube prototype: map (1), animation (2), transition to space–time cube (3), addition of set-time cube (4) and a net-time cube (5), transition to the superimposed view (6) and the juxtaposed view (7). Arrows indicate animated canvas transitions.

collection, in which she explained all visualization and interaction options. By a mixture of guidance and onboarding techniques, we aimed to transition from familiar to unfamiliar views, increase the visual complexity stepwise, and minimize coherence breaks. Figure 6 gives an overview on the guided walk-through: We started with a conventional map (1) and showed the temporal development of the collection with an animation (2). Then, a seamless canvas transition guided the user to the geographic STC representation (3). This view was amended stepwise with a categorical (4) and a relational STC (5). Afterward, the combined STCs were transformed to a SI (6) and a JP view (7) by means of seamless canvas transitions. After each step, users described their understanding of the visualization and any misconceptions were cleared up immediately. In addition the users were asked to compare the different views, to describe perceived strengths and weaknesses, and to state their overall preferences.

In phase II, the test participants could freely explore the IMDB movie collection with the same prototype. During the exploration, they were asked to think aloud and to share all observations on the data as well as on the prototype. During the whole procedure, an observer took notes of the participant's utterances and interactions. These protocols were analyzed qualitatively.

Results

Animated canvas transitions: After a seamless transition from a 2-D map to a geographic STC, eight users correctly understood the visualization technique (P2: "pictures are distributed according to time and space"), though three of them were not completely confident. One user required an explanation and another one mistook the yearly geographic dispersion for the overall number of pictures in this year.

The animated transition from the combined STC to the JP view helped all users to understand it quite easily (P2: "superimposed view, but now split up by time"). In contrast, the animated transition from the combined STC to the SI view was more difficult for the users: Only five users could correctly describe the new view (P10: "the time information was now extracted from the 3-D view, the color shows the chronology"), but some needed to see the transition multiple times. Another four users had problems understanding the new color-coding and frequently mistook it for the categories.

When asked about their opinion on the canvas transitions, half of the users regarded the *transitions* between the views as helpful for their understanding. Three users found it only partly useful (P8: "it should be slower to be useful") and two found it not useful at all (P7: "it is a nice effect, but not helpful").

Overall, the results of the evaluation confirm that animated canvas transitions helped most users to understand the STC and the JP view. The SI view turned out to be rather difficult due to a coherence break of the data points' color-coding: The transition from categorical to temporal color coding was not included in the animated transition and led to several difficulties and misconceptions. One user suggested to keep the categorical color coding in SI and use additionally color saturation, for example, to encode time.

Preferences: During the walk-through, users were asked to state their *opinion on each syn-diachronic view*, as compared to the others (see Figure 7) and to describe their individual pros and cons. Nine out of ten users found the STC view best (P1: "offers always a reference point for orientation"/P3: "good overview on the map and the sets"/P8: "most informative"/P10: "cool, good for exploration, more substance"). The SI view was rated high regarding aesthetics, but low for its information value (P4: "aesthetically more pleasing and extraordinary, but for

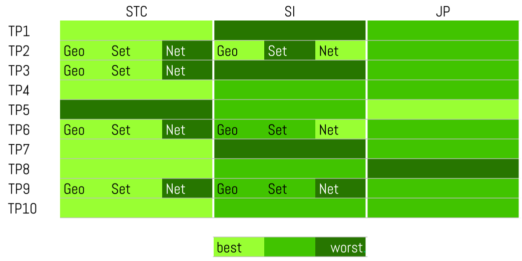

Figure 7. Preferences of different views.

analysis the 3-D view is better"). One exception was the network view, which two users found more comprehensible in SI (P2: "The 2-D network is more pleasant"). The JP view received mixed opinions. Most users found it useful for an analysis of single time points, but too scattered for an overall picture (P3: "it's too much, but good if you are interested in a year only"/P6: "if you want to look at a year only, it is useful, but for an overview other views are better"). ANI was found useful as a starting point (P2: "it's good in the beginning to see the genesis").

Synoptic insights: During the free exploration of the IMDB collection, we asked the users to think aloud and coded *multidimensional insights* into the data: Eight users had geotemporal insights (P2: "The U.S. is active over all years, in Europe there is a drop during WW2, and Asia started later only") and eight users reported a set-temporal insight (P10: "It is astonishing how early Sci-Fi films started."). Only one user had relational-temporal insights, whereas the other users had no insights at all in the network visualization or identified only influential films or clusters without a reflection on time.

Beyond these single views, we were interested to see how many users integrated more data dimensions in their insights, for example, sets, time, and space. We found such higher dimensional insights only in the think-aloud protocols of three test users (P3: "In the early 20th century, mainly dramas were produced in Europe."). An important and frequently used interaction technique for higher dimensional insights (as well as the reduction of information density and complexity) was the option to filter all views according to the categories, which was applied by seven test users.

These findings indicate on one hand that the integrated encoding could successfully support casual users in gaining synoptic insights into geographic space or categories over time. However, we did observe hardly any network-related insights over time. This might be due to the fact that the network view turned out to be the most complex view for most users and some even expressed difficulties to image its usefulness in general (P9: "It might have some logic").

On the other hand, our results do not indicate that the coordination of views was perceived as an effective coherence technique in supporting synoptic insights. While the juxtaposed and coordinated cubes or canvases of the system technically allow for a "massive parallel glimpse" on all time-oriented data dimensions, our users seemingly did not utilize this synoptic option during free exploration. Even though we observed some higher dimensional insights, these were gained by applying a categorical filter on the geotemporal view rather than by integrating patterns found in the set- and in the geotemporal view. However, in our evaluation procedure we did not explicitly prompt the users to look for synoptic patterns integrating all three perspectives. Given time, motivation, or more specific tasks, such insights might still occur.

DISCUSSION

The outlined system design contains numerous challenges for future work, and various aspects, which deserve further reflection.

Benefits and Costs of Coherence Techniques

From a cognitive science perspective, coherence is an important factor influencing the quality of internal representations—especially for representations of complex information spaces—and how effectively reasoning operations across multiple data dimensions can be performed. The PolyCube system implements three different visual coherence techniques, aiming to draw together different aspects of a multiview system (see Figure 2). Regarding the actual performance of these techniques, the results of our qualitative user study strengthen the notion that *animated canvas transitions* successfully act as an onboarding technique, which helped users to develop an understanding of a novel view coming from an already familiar one. Our results suggest that the animated transitions lead users to seamlessly integrate novel views with the existing internal representations of familiar views. The *integrated encoding* of time and geographic space as well as time and categories enabled the study participants to gain synoptic insights into the collection—an indicator of the integration of those data dimensions in the internal representation as well. However, the juxtaposition and *coordination of views* seemed to be less effective, as hardly any user connected geographic, categorical, and relational patterns. If they did, they used the categorical color-coding and filtering within the geotemporal view—a further argument for integrated encoding. In addition, *inconsistent encoding of visual variables (in our case: color)* was confirmed to be harmful for an overall understanding of a multiview system.[9] As this study is qualitative, we did not systematically investigate the effects of the different coherence techniques. Therefore, further research is required to assemble further visual coherence techniques and explore their benefits and limitations in a more systematic way.

While a strong motivation of this project emerged from the fact that the cognitive costs of multiple views can accumulate into solid barriers for casual users, we are aware of various overheads, which come along with coherent visualization design as well. Obviously, additional efforts regarding a system's design, implementation, and testing are needed to better connect and combine a plurality of views. Future studies have to corroborate our guiding hypothesis, and the beneficial effects of coherence techniques will have to be proven in a more systematic fashion, including novel evaluation methods.[7]

Utilizing Three Dimensions

Three-dimensional designs are frequently criticized in an information visualization context: Among other reasons, 3-D is known to generate occlusion of objects, depth ambiguity, perspective distortion of distances and angles, and to raise interaction costs for analysts trying to compensate these effects.[18] However, the resulting standard recommendation (i.e., to avoid 3-D and use multiple views instead) also incurs significant collateral costs. These costs (such as split attention and steep increase of cognitive load during postperceptual syntheses) frequently remain hidden during evaluation procedures, which focus on local and rather short-term analytical tasks, and which do not reflect on the role and relevance of information integration.[7] By contrast, the use of STC representations in the PolyCube system actually serves multiple integrative causes, and has additional positive effects, which arguably are also rarely heeded in the visualization context[5]:

- In line with SI, JP, and ANI, STC representations provide a visually integrated syn-diachronic view (coherence technique No. 2), thus helping to avoid the additional visualization of a linked timeline. However, STCs do so in a uniquely just and balanced manner: They utilize the most effective visual variable (i.e., positional encoding) equally for both sides: x- and y-axis to a synchronic, and the z-axis to the diachronic data dimension.[16]
- The morphology of the resulting point clouds can be parsed and read by the highly trained faculties of 3-D gestalt perception: Temporal developments become visible as spatial patterns along the vertical dimension—as effectively illustrated by the temporal pattern language encoded into the hull structure of the set-time cube. Empirical studies on casual users confirm that they can identify multidimensional patterns more quickly and more accurately with STC than with 2-D visualizations[17]. Similarly, the present study observed that they facilitate synoptic insights.
- As an additional feature, STC representations can act as translational hubs: They can mediate between the various syn-diachronic views (SI, JP, ANI), and also translate from temporal to nontemporal perspectives, while supporting visual analysts navigation by animated canvas transitions[11] (coherence technique No. 3).
- Studies show that STC representations have a certain attraction power and are considered "cool" (as summarized in reference[17]). Also in the present study, most casual users preferred the STC view for the exploration of a cultural collection. This kind of affinity should not be dismissed, given the importance of drawing casual users into an in-depth exploration process.

- Finally, as parallel, coordinated "PolyCubes" (coherence technique No.1), STCs provide a spatially adjacent "macro-shape" (so to say an "elephant"). This macroshape can also help us to navigate multidimensional datasets, and to transform abstract analytical tasks into movements of perspective-taking or of spatial manipulation (cutting, flattening, probing, comparing, etc.) of one or all cubes.[11]

Drawing these arguments together, we consider STC representations to provide a powerful and largely untapped potential for visual synthesis and information integration. By offering them as one out of four syn-diachronic views—reconciled and mediated by STC operations—we do not risk possible insights, but aim to maximize and connect them.

Flexibility

An important aspect in the development of PolyCube was the flexibility of the system. Depending on the data, the users, and their intentions, the system can be flexibly adapted: 1) The most relevant and interesting synchronic views (map, sets, or network) can be activated or deactivated, as can be additional visual encodings (such as the colors of data points, links, or hull structures). This allows users—who are easily overwhelmed by visual complexity—to reduce the information load to an appropriate level. 2) Different syn-diachronic views (STC, SI, JP, ANI) can be selected for exploration. This is especially important for casual users, who differ in their individual preferences (see Figure 7) and are more likely to explore a collection with views that meet their preferences. 3) Finally, PolyCube is intended to be open for all kinds of cultural collection data, which include temporal and geographic, categorical, or relational data. To enable the visualization of novel cultural datasets, the prototype pulls the data from a Google spreadsheet.

Missing and Uncertain Data

Data from historical collections and archives frequently contain significant gaps, ambiguities, and high amounts of uncertainty. For the many issues connected to such issues of heterogeneous data quality, the PolyCube prototype currently offers two coping techniques: 1) For collections with missing data dimensions (i.e., geoinformation or relations between objects), the number of views can be reduced (e.g., set-time-visualizations only, or geo-and-set-time combination only). 2) For uncertain data, we have started to develop glyph-based and relation-based uncertainty visualization techniques, which is a nontrivial endeavor, when consistent encoding is sought for an already complex visualization system.[19]

Limitations

Due to its computing-intensive character, we consider the STC view to be the system's bottleneck perspective—especially with regard to the network view, where large numbers of edges have to be visualized in addition to each object. While our two use cases (Cushman and IMDB) feature about 2000 objects per cube, we have been able to represent up to 20 000 objects per cube, while maintaining reasonable response times. However, we expect that analysts have to avoid the 3-D view, when they want to investigate large collections, going significantly beyond this threshold.

Generalization and Future Work

What we stated at the beginning about cultural collection data—rich, complex, and dynamic—actually applies for a wealth of other phenomena, not only, but especially in the (computational) social sciences and (digital) humanities realm. Novel visualization approaches and collaborations for these challenging subject matters are developing[20]—and the availability of multiple "interpretive" perspectives is a constant demand, as indicated by terms like *plurality, generosity, triangulation, and parallax*.[5] From event-based data (historical events) to trajectory data (biographies, military campaigns), relational data (historical networks), and infinite combinations of contextual visualizations of digital history, humanities, and cultural heritage studies—new types of visualization systems and visual–analytical environments are needed in the digital humanities. As an example, the PolyCube system concept has already been extended to represent biographical data in multiple information spaces, including nongeographic domains.[16]

CONCLUSION

With this article, we introduced the PolyCube visualization system for rich, time-oriented collection data with the specific aim to maximize synergies and connections between the encoding and representations of multiple data dimensions. While firmly embracing the use of multiple views, the system also aims to facilitate synoptic insights and cross-dimensional reasoning with less cognitive effort. For that matter, our main design strategies in the development of the prototype were 1) assembly and coordination of complementary visualizations, 2) integrated encoding of time with multiple types of syn-diachronic views, as well as 3) mediating the switching between various views with animated canvas transitions. An evaluation with casual users confirmed that especially the latter two are efficient techniques to create more integrated internal representations, while future work has to shed light on cognitive top-level syntheses with more systematic study designs.

Concerning the future application of the outlined system design, we consider an infinite amount of subject matters in digital humanities and social science domains to be connected to complex, dynamic systems, or time-oriented phenomena. To convey research and teaching communities with appropriate visualizations, the provision of many parallel views is not enough. With this article, we make the case for the development of more considerate "visual–synthetical" design strategies. While obviously providing all relevant visual–analytical functions, novel design strategies should also help users to integrate and connect multiple perspectives to complex mental models on a visual top-level of reasoning operations. With one possible approach to this kind of interface design, we aim to give an impulse to corresponding developments and discussions in a variety of related visualization domains.

ACKNOWLEDGMENTS

The authors would like to thank Emőke Ágnes Horvát for sharing the data of the study on cinematic citations.[12] This work was supported by a grant from the Austrian Science Fund (FWF, Project P28363-24), and the NFB (SC16-032).

REFERENCES

1. S. Bitgood, "Museum fatigue: A critical review," *Visitor Studies*, vol. 12, no. 2, pp. 93–111, 2009.
2. F. Windhager *et al.*, "Visualization of cultural heritage collection data: State of the art and future challenges," *IEEE Trans. Vis. Comput. Graphics*, vol. 25, no. 6, pp. 2311–2330, Jun. 2019.
3. M. Dörk, C. Pietsch, and G. Credico, "One view is not enough. High-level visualizations of a large cultural collection," *Inf. Des. J.*, vol. 23, no. 1, pp. 39–47, 2017
4. J. C. Roberts *et al.*, "Multiple views: Different meanings and collocated words," *Comput. Graph. Forum*, vol. 38, no. 3, pp. 79–93, 2019
5. F. Windhager, S. Salisu, G. Schreder, and E. Mayr, "Orchestrating overviews: A synoptic approach to the visualization of cultural collections. remaking collections," *Open Library Humanities, SI Remaking Collections*, 2018.
6. G. Schreder, F. Windhager, M. Smuc, and E. Mayr, "A mental models perspective on designing information visualizations for political communication," *J. eDemocracy Open Govern.*, vol. 8, no. 3, pp. 80–99, 2016
7. F. Windhager, G. Schreder, and E. Mayr, "Designing for a bigger picture: Towards a macrosyntax for information visualizations," *OSF Preprint*, Mar. 31, 2017, doi:10.31219/osf.io/q9kdt.
8. N. Kerracher, J. Kennedy, and K. Chalmers, "The design space of temporal graph visualisation," in *Proc. Eurograph. Conf. Visualization*, 2014, p. 1–5.
9. M. Q. W. Baldonado, A. Woodruff, and A. Kuchinsky, "Guidelines for using multiple views in information visualization," in *Proc. Conf. Adv. Vis. Interfaces*, 2000, pp. 110–119.
10. S. McCloud, *Understanding Comics: The Invisible Art.* New York, NY, USA: Harper Collins, 1994, pp. 60–93.
11. B. Bach, B. Dragicevic, D. Archambault, C. Hurter, and S. Carpendale, "A descriptive framework for temporal data visualizations based on generalized space–time cubes," *Comput. Graph. Forum*, vol. 36, no. 6, pp. 36–61, 2017.
12. A. Spitz and E. Á. Horvát, "Measuring long-term impact based on network centrality: Unraveling cinematic citations," *PloS One*, vol. 9, no. 10, 2014, Art. no. e108857.
13. M. Brehmer, T. Munzner, "A multi-level typology of abstract visualization tasks," *Trans. Vis. Comput. Graphics*, vol. 19, no. 12, pp. 2376–2385, Dec. 2013.

14. S. Salisu, E. Mayr, V. Filipov, R. Leite, S. Miksch, and F. Windhager, "Shapes of time: Visualizing set changes over time in cultural heritage collections," in *Proc. Poster Comp. EuroVis*, pp. 45–47, 2019.
15. S. Kriglstein, M. Pohl, and M. Smuc, "Pep up your time machine," in *Handbook of Human Centric Visualization*, New York, NY USA: Springer, pp. 203–225, 2014.
16. E. Mayr and F. Windhager, "Once upon a Spacetime: Visual storytelling in cognitive and geotemporal information spaces," *SPRS Int. J. Geo-Inf.*, vol. 7, no. 3, pp. 1–18, 2018.
17. E. Mayr, G. Schreder, S. Salisu, and F. Windhager, "Integrated visualization of space and time: A distributed cognition perspective," *OSF Preprint*. Dec. 2018, doi: 10.31219/osf.io/agvhw
18. M. Sedlmair, T. Munzner, and M. Tory, "Empirical guidance on scatterplot and dimension reduction technique choices," *IEEE Trans. Vis. Comput Graphics*, vol. 19, no. 12, pp. 2634–2643, Dec. 2013.
19. F. Windhager, S. Salisu, and E. Mayr, "Exhibiting uncertainty: Visualizing data quality indicators for cultural collections," *MDPI Informat.*, vol. 6, no. 3, pp. 1–16, 2019.
20. A. J. Bradley *et al.*, "Visualization and the digital humanities: Moving toward stronger collaborations," *IEEE Comput. Graph. Appl.*, vol. 38. no. 6, pp. 26–38, Nov./Dec. 2018.

Florian Windhager is a Research Associate with the Department for Knowledge and Communication Management, Danube University Krems, Austria. His research interests include information visualization, digital humanities, and new fields of application for visual reasoning. He received a Master of Philosophy degree and the Ph.D. degree in developing a visual analytics framework for the area of art-historical and biographical data. He is a Student Member of the IEEE. He is the corresponding author of this article. Contact him at florian.windhager@donau-uni.ac.at.

Saminu Salisu is a Research Associate with the Department for Knowledge and Communication Management, Danube University Krems, Austria, and a Ph.D. student with Vienna University of Technology, Austria. His research interests include information visualization, visual analytics, uncertainty visualization, and spatio-temporal visualization. Contact him at s.salisu@live.co.uk.

Roger A. Leite is currently working toward the Ph.D. degree with the Institute of Visual Computing and Human-Centered Technology, Vienna University of Technology (TU Wien), Vienna, Austria. His main research interests include information visualization, visual analytics, time-oriented problems, and networks. Contact him at roger.leite@tuwien.ac.at.

Velitchko Filipov is a Ph.D. student with the Institute of Visual Computing and Human-Centered Technology, Vienna University of Technology (TU Wien), Vienna, Austria. His research interests include information visualization, visual analytics, graphs, and networks. He is a Graduate Student Member of the IEEE. Contact him at velitchko.filipov@tuwien.ac.at.

Silvia Miksch is a Professor with the Institute of Visual Computing and Human-Centered Technology, TU Wien, Vienna, Austria. She served as paper cochair of IEEE VAST and EuroVis conferences and on the editorial board of IEEE TRANSACTIONS ON VISUALIZATION AND COMPUTER GRAPHICS. Her main research interests include visualization (particularly focus+context and interaction) and time. Contact her at miksch@ifs.tuwien.ac.at.

Günther Schreder is a Research Associate with the Department for Knowledge and Communication Management, Danube University Krems, Austria. His research focuses on how technology can support the cognitive processes of individuals and groups. He received a Master of Science degree in psychology from the University of Vienna, Austria. Contact him at guenther.schreder@donau-uni.ac.at.

Eva Mayr is a Postdoctoral Researcher with the Department for Knowledge and Communication Management, Danube University Krems, Austria. Her main research interests include the cognitive processes during interaction with information visualizations, in particular in "casual," informal learning settings. She received the Ph.D. degree in applied cognitive and media psychology from the University of Tübingen, Germany. Contact her at eva.mayr@donau-uni.ac.at.