Vision, Modeling, and Visualization (2020)
J. Krüger, M. Niessner, and J. Stückler (Eds.)

Visual Exploration of Cultural Heritage Collections with Linked Spatiotemporal, Shape and Metadata Views

S. Lengauer[1], A. Komar[1], S. Karl[2], E. Trinkl[2], R. Preiner[1], T. Schreck[1]

[1]Institute of Computer Graphics and Knowledge Visualisation, Graz University of Technology
[2]Institute of Classics, University of Graz

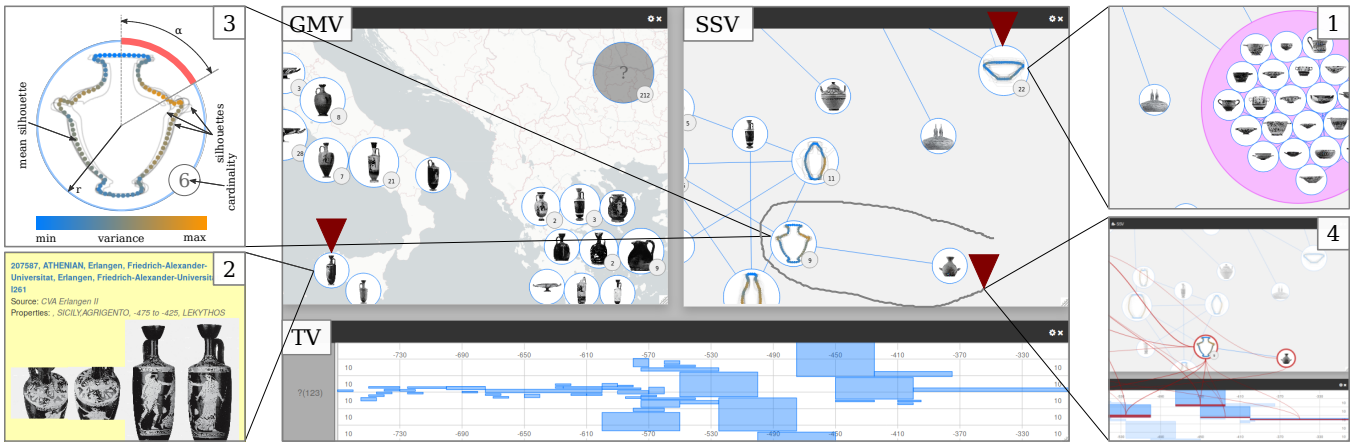

Figure 1: *Our Linked Views Visual Exploration System (LVVES) allows to explore space (GMV), time (TV) and shape (SSV) modality in an integrated system with linked views. In each view the CH artefacts are aggregated and clustered with regard to the respective modality and a dynamic degree of visual granularity, from overview (1) to close-up (2). We introduce multi-object previews (3) tailored to the SSV, providing visual cues regarding the artefacts substituted by an aggregation. Intra-view cluster correlations are visualised with alluvial-flow diagrams (4) spanning between views. The selection for this display happens twofoldly: (i) manual selection by the user with a lasso tool or (ii) automatically determined selection based on a cross-correlation measure between different layouts. Red arrows indicate mouse interactions.*

Abstract

The analysis of Cultural Heritage (CH) artefacts is an important task in the Digital Humanities. Increasingly, rich CH artefact data comprising metadata of different modalities becomes available in digital libraries and research data repositories. However, the large amounts and heterogeneity of artefacts in these repositories compromise their accessibility for common domain analysis tasks, as domain researchers lack a structural overview of the spatial, temporal, and categorical traits of the artefacts in these collections. Still, researchers need to compare artefacts along different modalities, put them into context, and deal with possible uncertainties, subjectivities, or missing data. To date, many works support domain research via interactive visualisation. The majority relies primarily on visualisation of text and metadata including spatiotemporal, image and shape data. However, fewer consider these types of data in a tightly coupled way. We present an approach for tightly integrated multimodal visual exploration of large CH data collections along space, time and shape traits. Based on requirements obtained in collaboration with domain researchers, we introduce a set of interlinked views for exploration of said modalities. An appropriately defined approach automatically computes most significant correlations across different modalities, guiding the user towards detecting interesting artefact relationships. We apply our approach to pertinent archaeological data collections, and demonstrate that characteristic explorative tasks are effectively supported and domain-relevant artefact relations can be discovered.

CCS Concepts

• ***Information systems*** → ***Digital libraries and archives;*** • ***Human-centered computing*** → ***Visualization; Visualization techniques; Interactive systems and tools;***

DOI: 10.2312/vmv.20201196

1. Introduction

The analysis of Cultural Heritage (CH) artefacts in terms of their historical context is of fundamental importance for archaeological research. To this end, finding and understanding relations between artefacts regarding traits of different inherent and derived modalities, e.g., spatial (findspot, provenance), temporal (dating), or categorical (shape, material, style), is a crucial task that often stands at the beginning of gaining new insights relevant to the domain. Thanks to the effort of Digital Humanities, more and more artefacts are digitized and collected in databases together with the relevant metadata that encode these traits and allow for efficient queries for relevant comparable artefacts using classic 'search by metadata'. While this supports the search for artefacts with given specific traits input by the user, important insights for the domain experts are often gained by analysing large volumes of artefacts as a whole, e.g., by investigating the relations between clusters within different trait modalities. However, classic search-based interfaces to digital repositories hardly allow archaeologists to gain a suitable overview of the whole volume of artefacts present in a collection, rendering the exploration and analysis of potentially interesting relations of artefacts across different modalities a tedious process. Moreover, in practice digital repositories typically contain artefacts with missing, unknown or uncertain metadata, additionally challenging the assessment of connections between groups of artefacts based on linked modalities. Finally, much of the information relevant for assessing the relations between CH artefacts are contained in artefact-related traits like their shape, painting style, or similar, which are mostly captured in images, and are rarely encoded in a low-dimensional metadata domain to search in.

In this work, we utilize the potential of digital collections for a more task-oriented analysis and exploration of CH data. We propose a visual analysis and exploration approach, that addresses the above-mentioned challenges by supporting the domain user in the analysis of large digital artefact collections and the assessment of significant relations between comprised artefacts based on multimodal traits. We employ multiple linked views revealing the localization, distribution and concentration of artefacts within a spatial, temporal, and a shape-based domain, as well as their connections across these modalities. In particular, to provide an overview of artefact records based on their shape, we propose a novel interactive graph-based representation that structures artefacts based on a custom shape-similarity metric, while supporting adaptive levels of visual granularity by introducing multi-object previews tailored to shape aggregations. To this end, we generalize an established graph sparsification technique for locally varying link densities to work for fully connected, undirected and weighted graphs. Moreover, we introduce a new cross-correlation measure for graph layouts, that allows an automatic detection and highlighting of prominent correlations of artefact similarities across different modality domains, providing guidance to potentially interesting connections between relevant artefact subgroups. Views on different modalities naturally incorporate missing and uncertain data of individual artefacts by respective visual cues.

The proposed design concepts are tailored into an interactive visual exploration system for artefact collections that exhibit shape and spatiotemporal metadata. While not being limited to this domain, in this paper we focus on the CH domain. In particular, we analyze and evaluate our system at the example of pottery artefacts, which are of major importance within this domain, and are preserved from prehistoric time on in large quantities due to the relative durability of its used ceramic material. The ancient Greek pottery, commonly called *vases*, is the most outstanding group in archaeological pottery. Besides using it for dating archaeological contexts, these vases shed a light on religion, daily life and society of Greek and Greek-influenced civilizations in the Mediterranean due to their specific shapes and painted decorations.

We evaluate our system based on a meaningful set of tasks defined by domain experts. The potential of our approach is demonstrated at the example of Greek pottery collections from a real-world repository, revealing relevant relations about the historic context of ancient pottery in an intuitive and easily-accessible way.

2. Related Work

To date, researchers have investigated many concepts for combined visualisations for spatiotemporal data and relational data. For example, [CSB*20] shows sequences of networks obtained from moving entities via node-link diagrams as well as glyph representation. Novak et al. [NMM*14] apply graph layouts to show historical social network relations together with the temporal aspect in a separate timeline. Preiner et al. [PSK*20] combine these aspects into an integrated network visualisation on top of a height field displaying the temporal aspect. In [WSL*20] the space time attributes of CH artefacts are presented within a three-dimensional PolyCube system. Network visualisations have also been used for abstract CH artefact characteristics, like vessel shape similarity by Van der Maaten et al. [vdMBL*06] or for displaying similarities across Mayan inscriptions [BFPM18; RPOG11]. In [Cob16] a map view of artefacts is proposed which can reveal semantic relationships between artefacts via visual links. Multi-view approaches, as ours, have seen widespread use as noticeable from the survey by Roberts et al. [RAB*19]. Important previous works in multi-view and linked visualisations include the Vis-Trails framework [BCS*05], allowing to connect individual views together in an interactive environment. Furthermore, the Improvise framework [Wea04] was also influential in defining linked view arrangements for data exploration. A system which allows to investigate multiple modalities across different views is presented by Steed et al. [SGCT20]. In our work, we apply linked views relying on connectors, to integrate appropriate shape, temporal, and spatial information in a self-contained visual exploration system for heterogeneous CH data. Hence, we include a set of highly relevant data modalities in domain analysis workflows. A systematic overview of visualisation approaches for the CH domain is given in [WFS*18].

3. Domain Analysis Tasks

In discourse with domain researchers we gained insight into their research workflows, based on which we determined a set of tasks relevant for an archaeological exploration tool. The tasks can be categorized as being of single-modal (**T1**-**T4**) as well as cross-modal (**T5**-**T6**) type and include: **T1** Gaining an overview of the geographic distribution of the provenances of artefacts, **T2** Examining patterns in the temporal distribution of artefacts, **T3** Gaining

an overview of the diversity and variability of artefact shapes, **T4** Examining artefacts at a dynamic level of visual granularity, down to close-up, **T5** Revealing connections between spatial, temporal and shape-based distributions and **T6** Evaluating salient connections determined by a cross-correlation measure.

4. Linked Views Visual Exploration System (LVVES)

Our system design is tailored to the analysis tasks specified above. Roberts et al. [RAB*19] present an established means to visualise multiple traits of attributed data via a *multiple view* approach. Consequently, we decided to display each of our supported modalities (shape, time and provenance) in a separate view, permitting the user a systematic inter-modality exploration. In general, multiple views have the downside that the user's attention is focused at one view at the time, making it hard to grasp the overall structure of the data and spot intra-modality correlations. Common ways to overcome this issue are coordinating, synchronizing or linking actions between individual views. For our problem at hand the latter, also referred to as *multiple linked views* [RAB*19], proved to be the most convenient. In order to work with a huge amount of artefacts at the same time a degree of aggregation or clustering is required. We design the system such that individual artefacts are perceptible at the lowest level while the views become more and more abstract at higher levels. This is referred to as *visual granularity* by [WFS*18] and its appropriate level is an essential design choice for any CH InfoVis system. The LVVES supports a dynamic level of visual granularity, meaning that there is at all time approximately the same number of previews visible but the level of detail changes based on user navigation interaction. In the reminder of the section the individual views for our selected modalities are described in detail. All of them have an underlying graph data structure $G(V,E)$ in common, where the vertices V (common across views) correspond to the displayed artefacts A and weighted edges E can be used to model the similarity between between artefacts. These graphs serve as basis for a *force directed layout* [FR91], where the positioning of an artefact (visualised with a preview) in a view is dictated by a set of forces acting simultaneously upon the vertices.

4.1. Geographic Map Viewer (GMV)

Geographic maps are an established approach to visualise spatial traits of CH artefacts [WMS*16], as we require in task **T1**. An artefact's origin within our map is indicated by a *preview*. If the preview stands for a single artefact, it is referred to as *Single Object Preview* (SOP) [WFS*18] and can exhibit 3D models and images as well as text and metadata. As all artefacts from our data base have associated image data we decided to display one representative image, embedded in a circular marker, as preview (Fig. 1 (GMV)). If a preview is descriptive for a set of artefacts we speak of *Multi-Object Previews* (MOP) [WFS*18] and arrangements like lists, grids or mosaics are common representatives. In our visualisation we display one arbitrary chosen image from the respective set of artefacts in such a case. Assigning an artefact to a specific point on a map requires a latitude, longitude coordinate pair, which is not the case for the provenance attribute associated with our artefacts. But they are given in the form of place names with highly varying degree of geographic expansion. Obtaining a latitude longitude coordinate pair for such a place name involves a process called *forward geocoding* and is offered by a multitude of web services. Nonetheless, a place name cannot be pinpoint to one exact location on the map but only to an area. We use the bounding boxes of these areas, which can be intersecting each other or even be contained within one another due to the large areas associated with some place names, as positional constraints for the previews, meaning that a preview can roam freely within it's associated bounding box. To get exact locations, as required for drawing the previews on the map, we use a force directed layout with gravitation force $\vec{f}_g(v)$ pulling vertices towards the center of their bounding box with constant strength. As soon as a vertex lies within it's bounding box, this strength becomes 0. Still, many previews of artefacts stemming from the same or intersecting areas are being painted fully or partially overlapping each other. We address this issue twofoldly. Firstly, a collision force $\vec{f}_{col}$ ($\vec{f}_{col} \gg \vec{f}_g$) ensures that previews are separated enough to avoid overlapping. Secondly, we aggregated vertices into surrogates, standing for multiple artefacts, based on common place names, resulting in a much smaller number of vertices $V_{agg} < V$. The number of artefacts represented by a MOP $|A_{MOP}|$ (henceforth referred to as cardinality) is visualised with a text label (Fig. 1 (3)) as well as a slightly larger preview radius r amounting to

$$r(|A_{MOP}|) = k\log(|A_{MOP}|) + r_{SOP}, \tag{1}$$

with $k = 0.3$ as a scaling factor and r_{SOP} as the single artefact preview radius. The absence of spatial information for some artefacts prevents their placement on the map. Yet, they can be of interest for domain users. We consider them by means of a 'placeholder' preview in whom all such artefacts are aggregated into. It is marked by a preview displaying a question mark and is always placed in the top right corner of the GMV as it can be seen in Fig. 1 (GMV).

4.2. Timeline Viewer (TV)

Linked timelines, *animation*, *superimposition* and *space-time cubes* are researched methods for displaying temporal information [WMS*16]. We decided for a timeline approach, as it is the only one visualising exclusively temporal information from the aforementioned, and we want to display exactly one modality per view. Aside from a mere timeline we also want to incorporate cardinality information for a certain date into our viewer. We implement this by means of a two-dimensional viewer with the temporal information encoded along x-dimension and the cardinality information encoded along the y-dimension. Yet, CH artefacts cannot be attributed to an exact point in time but only to a (possibly half-open) date interval due to uncertainties in the dating. Aggregating the artefacts by common time interval leaves us with a set of blocks we can place in our viewer. Although the expansion in x and y dimension of such a block is unambiguous the problem that many blocks are overlapping each other, due to overlappings in the corresponding date intervals, remains. This can be counteracted by rearranging the blocks along the y dimension such that they do not intersect each other. Finding an global arrangement which makes the optimal use of the available space is a variation of the *bin packing problem* and is combinatorial NP-hard [BV08]. However, we found for practical use an approximate solution to be sufficient. To this end, we apply a force layout with two counteracting forces: (i) a vertical force $\vec{f}_y$ governs the vertices to align around a vertical center line and (ii)

a larger collision force $\vec{f}_{col}$, with $\vec{f}_{col} \gg \vec{f}_y$, pushing blocks apart which are intersecting each other. Fig. 1 (TV) shows one possible result of this layout concept. Note that the TValso features a labelled grid facilitating the orientation. Same as with the GMV, we face the problem that many artefacts are missing a temporal attribution. We account for those with a designated block in a secluded area to the left of the timeline.

4.3. Shape Similarity Viewer (SSV)

The SSV provides an abstract visualisation of the shape variability and variability present in a collection. As in [vdMBL*06], we assume that an artefact's shape characteristics can be derived from its silhouette. To this end, we make use of the artefacts' associated image data, which is subjected to a manual filtering to discard image data showing only close-ups on surface details or related drawings. Subsequently, the images $I(a)$ for an artefact $a \in A$ are pre-processed as described by [LKL*20], before the largest contour of each image is extracted with the algorithm by [Suz*85]. To quantify the similarity between two contours we need a suitable representation of them, referred to as *feature descriptor*. We selected one state-of-the-art feature descriptor, the *Shape Contour Descriptor* (SCD) by Attalla and Siy [AS05], which is based on three different characteristics of a contour, e.g., the normalized distance of a segment from the center. With $c_{max}(i)$ as the largest contour in image i and $\mathbf{f}_{scd}(c_{max}(i))$ as the feature vector for i, $F(a) = \{\mathbf{f}_{scd}(c_{max}(I(a)_i) : 1 \leq i \leq |I(a)|\}$ denotes the set of feature vectors for artefact a. We define the similarity d between two artefacts a_i and a_j as the minimum Euclidean L_2-distance $d(a_i,a_j) = \min_{\{\mathbf{f}_k,\mathbf{f}_l\} \in X} L_2(\mathbf{f}_k,\mathbf{f}_l)$ between all feature vector pair combinations $X = \{\{F(a_i)_k, F(a_j)_l\} : 1 \leq k \leq |F(a_i)|, 1 \leq l \leq |F(a_j)|\}$.

The similarity relations between all artefacts are given by the symmetric square adjacency matrix $\mathbf{B} \in \mathbb{R}^{|A|\times|A|}$ with $b_{i,j} = b_{j,i} = d(a_i,a_j)$, which can be treated as a finite weighted graph $G^{SSV}(V,E)$ with vertices $V \mathrel{\widehat{=}} A$ and the edges E given by $E = \{\{i,j,w_{i,j}\} : 1 \leq i < j \leq |V|, w_{i,j} \propto b_{i,j}\}$. In the SSV, the grouping of artefacts based on their mutual shape similarity is done by a force layout where a link force $\vec{f}_l$ between all possible pairs of artefacts, e.g., $a_i, a_j \in A$ is proportional to the corresponding edge weight $\vec{f}_l(i,j) \propto w_{i,j}^{-1}$, assuring that similar artefacts are positioned close together while dissimilar ones are forced to stay further apart. Three additional forces act simultaneously upon all vertices: (i) a collision force $\vec{f}_{col}$, (ii) a strong repulsion force $\vec{f}_{rep}$, with a strength decreasing exponentially with distance, meaning that vertices not connected by a strong link force are pushed apart, resulting in a generally 'cleaner' overall layout, and (iii) a counteracting weak gravitation force $\vec{f}_g$ ($\vec{f}_g \ll \vec{f}_{rep} < \vec{f}_l \ll \vec{f}_{col}$) pulling all vertices towards a common origin and preventing weekly linked groups from drifting off. Additionally to the previews at the corresponding vertex positions, links are drawn as straight lines (Fig. 1 (SSV)).

4.3.1. Graph Aggregation

With one vertex per artefact the resulting graph layout will be huge, hindering the perception of a collection on a macro level. We want to provide a global overview by means of abstraction as in the GMV and the TV, where it is done with aggregation by common provenance and common date respectively. In the case of the SSV we do not have a comparable finite set of common values as a basis for aggregation, but we combine vertices iteratively until a desired level of aggregation is reached. We implemented a greedy pairwise aggregation which removes vertices connected to the edge with the lowest edge weight $\hat{w} = \min_{\{i,j,w_{i,j}\} \in E_{agg}} w_{i,j}$ from the set of vertices and the corresponding edge from the set of edges. This step is conducted iteratively, with $V_{agg}^{(0)} = V$, $E_{agg}^{(0)} = E$ at the initial step and $V_{agg}^{(k)} = V_{agg}^{(k-1)} \setminus \{i,j\}$, $E_{agg}^{(k)} = E_{agg}^{(k-1)} \setminus \{\{i,j,w_{i,j}\} \in E_{agg}^{(k-1)} : w_{i,j} = \hat{w}_{i,j}\}$ at the k-th step, until $|V_{agg}|$ reaches a value equal to the mean number of previews in the GMV and the TV, to ascertain similar visual granularity across all views. The cardinality of an aggregated vertex is encoded into the preview size r_{MOP} (Eqn. (1)).

4.3.2. Graph Sparsification

Initiating the force layout with $\vec{f}_l > 0$ between all vertex pairs results in a hairball-like structure due to the layout being heavily overconstrained. The process of thinning out these link forces while preserving link-induced accumulations at the same time is generally referred to as *graph sparsification*. Due to similarity values between all vertices the aggregated graph G_{agg}^{SSV} features the highest possible number of edges with $|E_{agg}| = |V_{agg}|(|V_{agg}|-1)/2$, referred to as *complete* graph. The goal of sparsification is to find a subset of edges $E_{sparse} \subset E_{agg}$ with $|E_{sparse}| \ll |E_{agg}|$. The trivial approach is to remove a fixed portion of edges with the highest weights [SPR11], or omitting all edges exhibiting a weight above a globally defined threshold t_{glob}, yielding $E_{sparse} = \{\{i,j,w_{i,j}\} \in E_{agg} : w_{i,j} \leq t_{glob}\}$. A critical flaw of this approach is that the locally varying edge weights are not taken into account, resulting in the dissolution of some clusters while too many links are preserved in others.

Hence, we base our sparsification on an approach which aims at discarding inter-cluster edges while retaining intra-cluster edges. More specifically, we utilize the method described by Saturuli et al. [SPR11] (see Algorithm 2 in their work), which we adopt for complete graphs in the sense that we do not use the Jaccard similarity coefficient for determining the similarity between two vertices v_i, $v_j \in V$ but our own definition based on the distance of their neighboring edges given by

$$d_{nn}(E_i,E_j) = \frac{1}{N}\sum_{k=0}^{N} |k - indexOf(E_{i_k},E_j)|, \tag{2}$$

with E_i as all edges incident to v_i and E_j as all edges incident to v_j, N corresponds to the number of edges $N = |E_i| = |E_j|$ and $indexOf(e,E)$ returns the index of edge e in a set of edges E. The set of incident edges $inc(E,v)$ for a vertex $v \in V$ and edges E is given by

$$inc(E,v) = \{\{i,j,w_{i,j}\} \in E : v = i \vee v = j\}. \tag{3}$$

The edge in E_{agg} ($E_{agg}^{(0)} = E_{sparse}$) with the greatest distance $\hat{e} = \max_{\{i,j,w_{i,j} \in E_{agg}\}} d_{nn}(inc(E_{agg},i), inc(E_{agg},j))$ between it's connected vertices is removed iteratively until E_{agg}/E_{sparse} drops below a predefined threshold.

4.3.3. Shape Preview Surrogate (SPS)

Same as with the GMV we require MOPs for aggregated vertices. As the SSV deals exclusively with artefact silhouettes we con-

cluded that such previews should embed this information. To this end, we implemented a glyph which displays the Euclidean mean of a set of silhouettes together with a colour coding implying the local variances of silhouette shapes as depicted in Fig. 1 (3). For a set of aggregated artefacts A_{agg}, associated with a MOP, this involves normalization and registration of all the representative silhouettes $H(A_{agg}) = \{c_{max}(I(a)_0) : a \in A_{agg}\}$. Let $\mathbf{s}_i$ be the vector of n points, sampled at equidistant positions along the registered silhouette $h_i \in H$. A SPS is described by a point vector $\mathbf{s}_{SPS}$, initialized with the pointwise Euclidean mean of two silhouette samples $\mathbf{s}_1$ and $\mathbf{s}_2$ with $s^{(0)}_{SPS_i} = \frac{1}{2}s_{1_i} + \frac{1}{2}s_{2_i}, 0 < i < n$. All addition silhouettes are added on-line, meaning that the k-th silhouette vector $\mathbf{s}_k$ can be added with $s^{(k+1)}_{SPS_i} = s^{(k)}_{SPS_i} + k(s_{k_i} - s^{(k)}_{SPS_i}), 0 < i < n$. Additionally, we calculate the standard deviation of each point in $\mathbf{s}_{SPS}$. The glyph we create from this information consists of a series of dots representing the mean silhouette on a neutral background with the standard deviation encoded in the colour. Same as in the GMV the size of glyph in the viewer scales with its cardinality according to Eqn. (1).

5. Analytical and Interactive Support

The visual connection across views is established via user interaction, in the way that the user selects a subset of the data in one of the views by a selection action, which involves in many cases brushing, or, in our case, a selection lasso. This subset is then highlighted in all views by changing colour or size. We take it one step further and also draw links between views in order to visualise the correlation of clusters and aggregations across modalities (Sec. 5.2). Apart from this highlighting based on user selection, we also provide the user with an automatically calculated 'most significant' selection based on a cross-correlation measure (Sec. 5.3).

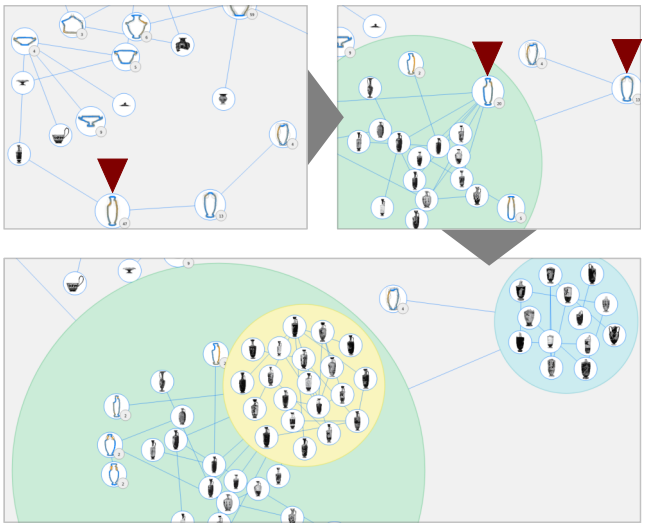

Figure 2: *Balloon views with different levels of recursion depth. The close-ups show joined vessels, all of the shape lekythos, in different groupings.*

5.1. Hierarchic Balloons

To allow a user to examine the data at all levels of visual granularity as required by **T4** we need a system for reversing parts of our aggregation dynamically. By now, this kind of exploration is is implemented for the SSV only, where we make use of the aggregated graph data structure $G^{SSV}(V_{agg}, E)$ (Sec. 4.3.1), where some of the vertices serve as surrogates for others but have a hierarchic list of vertices and edges, meaning that the data structure at hand constitutes a tree of graphs with the leafs corresponding to single artefacts $G_{leaf}(\{.\}, \emptyset)$. A method, suitable for the rest of our system, A proper visualisation technique for this kind of data is referred to by *balloon view* [HMM00]. The basic idea is to encapsulate subtrees in circles attached to the father node, isolating it visually while still preserving its relation to the next higher level in the hierarchy. In our implementation a balloon has its own force layout and is subjected to the same aggregation and sparsification operations (Sec. 4.3.1 and 4.3.2) as the graph of the root node, meaning that it can contain aggregated nodes as well. Those can also be turned into balloons recursively, down to the lowest level of the tree. Balloons are completely unaware of nodes higher up the hierarchy and incorporated balloons are considered only by means of their radius, determined by the extent of its force layout, which serves as parameter for the collision force. To avoid visual overloading, balloons are added only on demand, meaning that a user has to click its corresponding father node (Fig. 2). Note that a balloon is coloured with an arbitrary selected colour which is restricted to shades which have a sufficiently high contrast to the colour of its parent. It is possible to display multiple balloons from different levels of the hierarchy as well as from the same level.

5.2. Intra-view Highlighting and Linking

The linking of views in a multiple linked views system is referred to by *highlighting*. Typically, the user selected a subset of the displayed data by means of a brush or lasso [RAB*19] and subsequently the same artefacts in all views are highlighted by a change in colour, size or other factors. For our purpose a lasso selection turned out to be the easiest to use for choosing a selection. For highlighting purposes we decided for a change in colour, since the a preview's size already encodes cardinality information (Sec. 4.1). A challenge pose the different aggregations in different views, for oftentimes only a fraction of the artefacts related to a MOP belong to the current highlight. We visualise this partial highlight with a red arc around a MOP (Fig. 1 (3)), where the angle α amounts to $\alpha = |\text{"}num\ highlights\text{"}|/|\text{"}num\ artefacts\text{"}| \cdot 2\pi$. Due to the navigation capabilities of the views not all previews are within the currently displayed section all the time. If this is the case for a highlighted preview we want to indicate their presence regardless and do so via *signposts*.

Still, this form of highlighting does not allow to draw conclusions regarding the correlation of clusters across separate views. To this end we implement an alluvial-flow diagram, comparable to the diagram presented in [Sau17], between all views where the streams start and end at the respective previews, and the stream strength corresponds to the number of highlights (Fig. 5). A stream is only displayed if both, the source and the target preview are within the currently displayed sections.

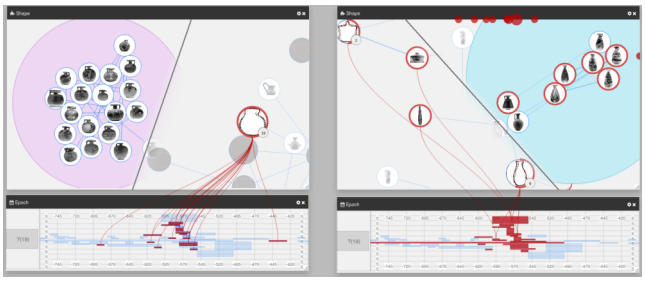

Figure 3: *Left: Selection of aryballoi SPS (individual artefacts in the pink balloon) to examine dating. Right: selection of same date range showing other shapes, e.g. - in the light blue balloon - the group of alabastra.*

5.3. Intra-view Correlation

Aside from a user-defined selection of intra-view links, we also pre-calculate the 'most significant' selections. In other words we determine those artefacts whose neighborhoods are the most similar across all views. Since we have graph data structures with a common set of vertices across all our views (G^{GMV}, G^{TV} and G^{SSV}) we do so by defining a vertex-wise correlation measure for different graph layouts. The intra-view correlation $ivc(v)$ for a vertex $v \in V$, common across graphs $G^i(V,E^i)$ and $G^j(V,E^j)$, is thus given by the similarity of its neighborhoods with $ivc(v) = d_{nn}(inc(E^i,v), inc(E^j,v))$ (Eqn. (2) and (3)). For n graphs and their set of edges $M = \{E^i : 1 \leq i \leq n\}$ the cross-correlation is given by

$$ivc(v) = \frac{1}{\binom{n}{2}} \sum_{\{E^0,E^1\} \in \binom{M}{2}} d_{nn}(inc(E^0,v), inc(E^1,v)). \quad (4)$$

For the SSV this measure is directly applicable as is has a complete graph $G^{SSV}(V,E^{SSV})$ defined by its adjacency matrix. The graphs of the TV and the GMV, on the the other hand, do not have any edges $G^{GMV} = G^{TV} = G(V,\emptyset)$, resulting in an uniform cross-correlation of 0. We compensate for this drawback by adding artificial edges between all vertices for G^{GMV} and G^{TV}. For the prior, the weight of edges between vertices belonging to the same cluster (and thereby same location) are set to 0, while all others are initialized with a weight relative to their geographic distance. For the latter, the weight of an edge between vertices belonging to the same date range are set to 0 while the weight of an edge between two vertices $v_i, v_j \in V$, belonging to different date ranges is defined as the difference $d_{TV}(v_i,v_j) = |z(v_i) - z(v_j)|$, where $z(v)$ indicates the center of the date interval associated with v. Note, that no normalization of edge weights is necessary for they serve only to determine an order. The resulting intra-view correlations are sorted by value and only the most prominent ones, down to a user defined threshold, are displayed via the highlighting and parallel links systems (Sec. 5.2), with the strength of the correlation encoded into the transparency of the links and associated previews (Fig. 4).

6. Results and Evaluation

To evaluate our approach we have implemented a prototype as a web application which relies heavily on the D3 JavaScript library in connection with HTML5 Canvas. A C++ back-end is responsible for the computation of the shape similarity required by the SSV (Sec. 4.3) as well as for the intra-view correlation (Sec. 5.3). The geocoding necessary for the GMV was obtained by the OpenCage Geocoder API. The LVVES has been used by domain experts, who also co-authored this publication, for the analysis of real-world archaeological data, following the different analysis tasks in Sec. 3.

6.1. Dataset

To test the LVVES we chose Greek vases as a very distinctive group of CH artefacts which are usually published on a standardized basis. We selected for this step of evaluation pictorial data as well as metadata. The sources comprise five printed fascicles of the Corpus Vasorum Antiquorum (CVA) – Dresden 2, in 2015; Erlangen 2, in 2007; Göttingen 3, in 2007; Jena 1, in 2011; and Munich 16, in 2010 – and the Beazley Archive Pottery Database (BAPD). The BAPD (https://www.beazley.ox.ac.uk/index.htm) is a freely accessible database for Greek pottery hosted in Oxford and collected 118,805 (accessed June 18, 2020) vases to date, from whom we composed a dataset for our showcases, encompassing 614 vases with 2,799 images. They have been filtered automatically to exclude fragments which do not exhibit meaningful shape information. From the remaining 317 artefacts the associated images have been filtered manually (Sec. 4.3). 61% of these vases have dating information, while only 32% have associated provenance information. This drawback is typically for museum collections of antiquities which are mostly compiled of artefacts coming from the art market since the 18th cent. AD. [SS14].

6.2. Application and Domain Expert Evaluation

To evaluate the system's mode of operation in archaeological field of research and to demonstrate some results we selected three specific showcases. For each of them we selected a different set of modalities (date, provenance and shape). For **T1** we use the GMV to investigate artefacts' provenances, for **T2** the TV to analyse the temporal distributions of artefacts and for **T3** the SSV to get an overview of shape variations.

For the first showcase, related to **T3**, **T4**, we take the data sets of Erlangen, Göttingen and Dresden and look at their representation in the SSV. Greek pottery was produced in different shapes and sizes according to the intended use (storage, drinking, containing perfume, etc.). Although the main shape types of Greek pottery remained fairly constant due to their function, there were a lot of varieties and differences in accentuation of details [Coo92]. E.g., one specific shape, the lekythos, in general a slender, single-handled and narrow-necked vessel, was used for storing oil. The SSV displays this common shape of lekythoi efficiently with a SPS of 47 artefacts (Fig. 2 top left). The SPS enables to recognise at a glance the low variability of the individual shape profiles of these clustered artefacts, as only small differences can be seen in the transition from the shoulder to the neck of these vessels. The dynamic level of visual granularity from overview to close-up by means of hierarchic balloons (Sec. 5.1) facilitates the examination of linked vases in detail. The first balloon (Fig. 2 top right) shows some varieties of lekythoi and a further extensive SPS of 20 artefacts. The second balloon (pink) demonstrates that SSV has clustered herein the

Figure 4: *The LVVES with intra-view links based on the most significant relations revealed by our intra-view correlation. Nolan amphorae and cups preferred in Etruria are detected. The two balloon views on the right side correspond to the highlighted SPSs in the SSV.*

lekythoi of the canonical proportions with narrow, almost cylindrical body. A limitation pose degenerated vases, e.g., different vase types with broken mouths are joined in one SPS in Fig. 2 bottom.

For the second showcase, addressing **T5**, we look at the vase collection of the university Jena in its entirety, more specifically on their represented shapes in relation to their given date ranges. Within the largest cluster, depicted in Fig. 3 left, the SSV grouped together vessels with a globular body. Apart from two exceptions - an early Protocorinthian aryballos and a late belly lekythos - all exhibit a shape known as *round aryballos*, a perfume vessel with a spherical body, a narrow neck and a broad disc-like mouth, which becomes regular at the end of the third quarter of the 7th cent. BC. in Corinth. This shape achieves a mass production with a wide distribution in the first half of the 6th cent. BC and disappears almost completely in the Greek sphere after the mid-6th cent. BC. The balloon in Fig. 3 left also shows imitations of this specific Corinthian shape in other vase production centres: in East Greek/Rhodes (made of Faience) and in Italy (the so-called Corinthianizing pottery). If we investigate which other shapes existing from the first half of the 6th cent. BC at Jena, the SSV offers, among some single shape types, a second larger group (Fig. 3 right). In this group all Corinthian and Corinthianizing alabastra at Jena were grouped together (with two outliers). To sum it up, the LVVES clearly supports the investigation of typo-chronological developments of vessel shapes and can enable the recognition of new strands of developments or peculiarities with a larger datasets.

The third showcase is a characteristic example for **T6**, connecting shape, provenance and date (Fig. 4). Neck amphorae in a distinctive slim version and with an elongated neck are clustered by the SSVwhile the GMV reveals that the vast majority of this type was found in Campania, Southern Italy, especially in the ancient city of Nola. This type of amphorae was named, following the archaeological site, *Nolan amphorae*. Even though this is not a new discovery, the visualisation emphasises export form Athens to Southern Italy. The ancient region of Etruria in Central Italy is a preferred exportation site for another vase shape, the cup. The timeline outlines that these exports started later and lasted longer than the exports of Nolan amphorae to Southern Italy. Looking at the pottery market in general, it is an interesting archaeological fact that specific shapes were popular in different region of Italy.

7. Limitations & Future Work

An open question is how to determine a sensible amount and routing for the parallel links in order to achieve the optimal balance between a maximum of visible information and a sensory overload, as it is the case in Fig. 5. To this end, we want to look at filtering approaches as well as methods routing with reduced collisions like the method presented by Steinberger et al. [SWS*11]. Another challenge are historic toponyms in artefacts' provenances and other spatial metadata which have no equivalent in modern geocoding services. We want to utilize a gazetteer which is specialized in ancient place names like the Pleiades project (https://pleiades.stoa.org/home). Extensions are also planed for the database as well as the number of supported modalities. The GMV can visualise any spatial data (e.g. production sites, current collections) while the SSV can display any non-categorical data (e.g. painting style) where relations between artefacts can be specified by a distance function. Lastly, we want to elaborate on our user guidance with systems capable of learning from user interaction for more customized suggestions based on collected information.

8. Conclusion

Gaining an insight and understanding on CH artefacts across modal boundaries is an archaeologist's everyday life. Available data is comprised of 3D scans, images, and a variety of metadata and helps

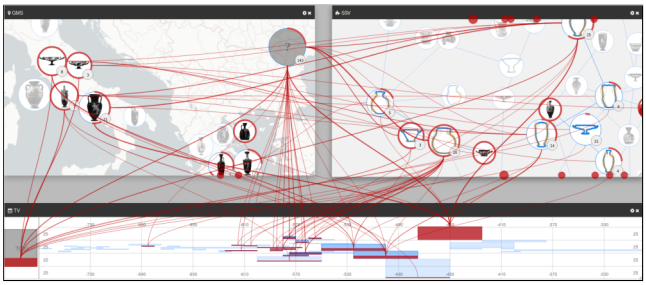

Figure 5: *Too many parallel links at the same time congest the LVVES and hinder the detection of interesting correlations.*

to analyze not only the economic situation of producers and recipients but the ancient social systems in general which is a core research question of the Classical Studies. The visual exploration system we present addresses these challenges with state-of-the-art visualisations for a selection of modalities and bring them together with visual links. The evaluation by domain experts has shown that our cross-modal visualisation and guidance enables expert users to work with a huge number of artefacts and establish connections between the artefacts' modalities. Even the first results for a small data set already reveals interesting and well-known relations. To sum it up, our prototype provides a bedrock for a system which can be used to handle state-of-the-art domain research questions.

Acknowledgments

This work was co-funded by the Austrian Science Fund FWF and the State of Styria (AT) by the project *Crossmodal Search and Visual Exploration of 3D Cultural Heritage Objects* (P31317-NBL).

References

[AS05] ATTALLA, EMAD and SIY, PEPE. "Robust shape similarity retrieval based on contour segmentation polygonal multiresolution and elastic matching". *Pattern Recognition* 38.12 (2005), 2229–2241 4.

[BCS*05] BAVOIL, LOUIS, CALLAHAN, STEVEN P., SCHEIDEGGER, CARLOS EDUARDO, et al. "VisTrails: Enabling Interactive Multiple-View Visualizations". *16th IEEE Visualization Conference, VIS 2005, Minneapolis, MN, USA, October 23-28, 2005*. 2005, 135–142 2.

[BFPM18] BOGACZ, BARTOSZ, FELDMANN, FELIX, PRAGER, CHRISTIAN, and MARA, HUBERT. "Visualizing Networks of Maya Glyphs by Clustering Subglyphs". (2018) 2.

[BV08] BERNHARD, KORTE and VYGEN, JENS. "Combinatorial optimization: Theory and algorithms". *Springer, Third Edition, 2005*. (2008) 3.

[Cob16] COBURN, JOHN. "I dont know what im looking for: Better understanding public usage and behaviours with tyne & wear archives & museums online collections". *MW2016: Museums and the Web* (2016) 2.

[Coo92] COOK, ROBERT MANUEL. *Greek Painted Pottery, Third edition*. Routledge, London, 1992 6.

[CSB*20] CAKMAK, EREN, SCHÄFER, HANNA, BUCHMÜLLER, JURI, et al. "MotionGlyphs: Visual Abstraction of Spatio-Temporal Networks in Collective Animal Behavior". *Computer Graphics Forum*. Vol. 39. 3. 2020 2.

[FR91] FRUCHTERMAN, THOMAS MJ and REINGOLD, EDWARD M. "Graph drawing by force-directed placement". *Software: Practice and experience* 21.11 (1991), 1129–1164 3.

[HMM00] HERMAN, IVAN, MELANÇON, GUY, and MARSHALL, M SCOTT. "Graph visualization and navigation in information visualization: A survey". *IEEE Transactions on visualization and computer graphics* 6.1 (2000), 24–43 5.

[LKL*20] LENGAUER, STEFAN, KOMAR, ALEXANDER, LABRADA, ARNIEL, et al. "A sketch-aided retrieval approach for incomplete 3D objects". *Computers & Graphics* 87 (2020), 111–122 4.

[NMM*14] NOVAK, JASMINKO, MICHEEL, ISABEL, MELENHORST, MARK, et al. "HistoGraph–A visualization tool for collaborative analysis of networks from historical social multimedia collections". *2014 18th International Conference on Information Visualisation*. IEEE. 2014, 241–250 2.

[PSK*20] PREINER, REINHOLD, SCHMIDT, JOHANNA, KRÖSL, KATHARINA, et al. "Augmenting Node-Link Diagrams with Topographic Attribute Maps". *Computer Graphics Forum* (2020). ISSN: 1467-8659 2.

[RAB*19] ROBERTS, JC, AL-MANEEA, HAYDER, BUTCHER, PWS, et al. "Multiple Views: different meanings and collocated words". *Computer Graphics Forum*. Vol. 38. 3. Wiley Online Library. 2019, 79–93 2, 3, 5.

[RPOG11] ROMAN-RANGEL, EDGAR, PALLAN, CARLOS, ODOBEZ, JEAN-MARC, and GATICA-PEREZ, DANIEL. "Analyzing ancient maya glyph collections with contextual shape descriptors". *International Journal of Computer Vision* 94.1 (2011), 101–117 2.

[Sau17] SAUNDERS, DAVID. "The distribution of the Berlin Painter's vases". *The Berlin Painter and His World: Athenian Vase-Painting in the Early Fifth Century B.C.* Ed. by PADGETT, J. MICHAEL. Princeton University Art Museum, 2017, 107–131 5.

[SGCT20] STEED, CHAD A, GOODALL, JOHN R, CHAE, JUNGHOON, and TROFIMOV, ARTEM. "CrossVis: A Visual Analytics System for Exploring Heterogeneous Multivariate Data with Applications to Materials and Climate Sciences". *Graphics and Visual Computing* (2020), 200013 2.

[SPR11] SATULURI, VENU, PARTHASARATHY, SRINIVASAN, and RUAN, YIYE. "Local graph sparsification for scalable clustering". *Proceedings of the 2011 ACM SIGMOD International Conference on Management of data*. 2011, 721–732 4.

[SS14] SCHMIDT, STEFAN and STEINHART, MATTHIAS. "Sammeln und Erforschen: griechische Vasen in neuzeitlichen Sammlungen". *Beihefte zum Corpus Vasorum Antiquorum Deutschland 6*. Verlag CH Beck, Munich, 2014 6.

[Suz*85] SUZUKI, SATOSHI et al. "Topological structural analysis of digitized binary images by border following". *Computer vision, graphics, and image processing* 30.1 (1985), 32–46 4.

[SWS*11] STEINBERGER, MARKUS, WALDNER, MANUELA, STREIT, MARC, et al. "Context-preserving visual links". *IEEE Transactions on Visualization and Computer Graphics* 17.12 (2011), 2249–2258 7.

[vdMBL*06] Van der MAATEN, LAURENS, BOON, PAUL, LANGE, GUUS, et al. "Computer vision and machine learning for archaeology". *Proceedings of Computer Applications and Quantitative Methods in Archaeology* (2006), 112–130 2, 4.

[Wea04] WEAVER, CHRIS. "Building Highly-Coordinated Visualizations in Improvise". *10th IEEE Symposium on Information Visualization (InfoVis 2004), 10-12 October 2004, Austin, TX, USA*. 2004, 159–166 2.

[WFS*18] WINDHAGER, FLORIAN, FEDERICO, PAOLO, SCHREDER, GÜNTHER, et al. "Visualization of cultural heritage collection data: State of the art and future challenges". *IEEE transactions on visualization and computer graphics* 25.6 (2018), 2311–2330 2, 3.

[WMS*16] WINDHAGER, FLORIAN, MAYR, EVA, SCHREDER, GÜNTHER, et al. "Reframing Cultural Heritage Collections in a Visualization Framework of Space-Time Cubes." *HistoInformatics@ DH*. 2016, 20–24 3.

[WSL*20] WINDHAGER, FLORIAN, SALISU, SAMINU, LEITE, ROGER A, et al. "Many Views Are Not Enough: Designing for Synoptic Insights in Cultural Collections". *IEEE Computer Graphics and Applications* 40.3 (2020), 58–71 2.