

histoGraph - A Visualization Tool for Collaborative Analysis of Networks from Historical Social Multimedia Collections

Jasminko Novak^{1,2}, Isabel Micheel¹

¹European Institute for Participatory Media

²University of Applied Sciences Stralsund
Berlin, Germany

j.novak@eipcm.org, i.micheel@eipcm.org

Lars Wieneke³, Marten Düring^{3,4}

³Centre Virtuel de la Connaissance sur l'Europe

⁴University of North Carolina at Chapel Hill
Sanem, Luxembourg

lars.wieneke@cvce.eu, marten.during@cvce.eu

Mark Melenhorst

TU Delft

Delft, Netherlands

m.s.melenhorst@tudelft.nl

Javier Garcia Morón

Homeria Open Solutions

Cáceres, Spain

javiergarcia@homeria.com

Chiara Pasini, Marco Tagliasacchi, Piero Fraternali

Politecnico di Milano

Milan, Italy

chiara.pasini@polimi.it, marco.tagliasacchi@polimi.it,

piero.fraternali@polimi.it

Abstract— This paper describes the design and development of histoGraph, an interactive tool for explorative visualization and collaborative investigation of historical social networks from multimedia collections. Developed in an interdisciplinary collaboration of computer scientists, historians, HCI researchers and interface designers, the tool aims at supporting historians in the discovery and historical analysis of relationships between people, places and events. A special focus is on the identification and interactive visualization of social relations from historical photo collections through a combination of automatic analysis and expert-based crowdsourcing. The tool design bridges the gap between established network analysis and visualization techniques and traditional hermeneutic research methods in historical research. It integrates visual exploration with hybrid social graph construction, hypothesis formulation and the consultation of digitized primary sources. A formative evaluation of the current prototype, developed as a domain-specific application for historians in the field of European integration points to opportunities and critical factors in applying this approach to support and further current research practices in digital humanities.

Keywords- Network Visualization, Digital Humanities, Knowledge Visualization, Information Visualization, Social Network Analysis, Collective Intelligence, Crowdsourcing, HCI, Visual Interfaces

I. INTRODUCTION

One major focus of current applications of information and knowledge visualization in the humanities is the analysis and discovery of patterns and relationships in large

information collections that are subject of humanist inquiry. A particular line of work deals with the application of computational methods to the identification and analysis of historical social networks and relationships (e.g. relationships between scholars and historical personalities [2][38][43], international organizations [42] or genealogical networks [22]). Such approaches typically focus on text corpora for identifying and analyzing social relationships.

At the same time, multimedia collections can offer rich insights and multimedia material is frequently used for "manual" inspection in areas of related research in the humanities (e.g. historical analysis of political processes such as European integration). This raises the question of what kind of relationships can be observed in and induced from such collections. How can the discovery of relationships between entities such as people, places and events from multimedia material be supported by visualization in a way that supports humanist inquiry?

Though a large body of work on underlying computational and visual techniques exists in the areas of multimedia information retrieval, knowledge discovery, visualization and HCI, its transposition to digital humanities is yet little investigated. There are several reasons for this; on the one hand, the application of techniques or tools for multimedia processing tends to be more complex than their counterparts from text-mining, requiring close interdisciplinary collaborations between researchers in the humanities and technical sciences. On the other hand, in the humanities there is a tradition of manually curating (semantically structuring) multimedia collections with great care - resulting in textual metadata that can be processed by text mining. The use of such textual metadata for analysis of

multimedia collections is thus based on a pre-selection that reflects a number of (partly implicit) choices that in effect predetermine the kind and range of possible research questions that can be asked and investigated or insights that can be gained.

Finally, the different epistemological traditions of the humanities and the technical disciplines make the transfer of techniques and inter-disciplinary collaboration all but straightforward. The humanities tend to view an often “objectivist” stance of the technical sciences as irreconcilable with their interpretivist approaches [11].

In this paper, we present an approach to addressing some of these issues in designing and developing a visual tool for explorative and collaborative investigation of historical social networks from multimedia collections. Developed in an interdisciplinary collaboration of computer scientists, historians, HCI researchers and interface designers, the tool aims at supporting historians in the discovery and historical analysis of relationships between people, places and events. Rather than applying a ready-made tool for network visualization and analysis, we consider and identify specific requirements of the process of humanist inquiry in this application domain. We discuss and propose specific design principles to which such tools should adhere and present a concrete realization of an application prototype turning them into practice.

In the current state of development, we place a special focus on the identification and interactive visualization of social relations from historical *photo collections* through a combination of automatic analysis and crowdsourcing. Beyond combining multimedia processing techniques with network analysis and visualization, the tool design aims at supporting fundamental principles of humanistic inquiry, such as interpretative research and constructivist approach to knowledge generation and visualization. To this end, it integrates visual exploration with collaborative social graph construction, hypothesis formulation and hermeneutic analysis of primary sources. A formative evaluation of the current prototype, developed as a domain-specific application for historians in the field of European integration points to the opportunities and critical factors in applying this approach to support and extend current research practice in digital humanities.

II. APPLICATION CONTEXT

The discovery of patterns and relationships in the primary sources by cross-referencing them is essential to humanities research and teaching. In particular, this includes the discovery and analysis of relationships between people, places and events in order to (re)define research hypotheses, identify correlations or establish causal relationships. Digital archives have quickly gained importance as resources for such research in the humanities. Frequently, textual collections are searched for occurrences of specific entities such as individuals or institutions in order to study their interaction with other entities: e.g. as in the analysis of the Smithsonian Cooper-Hewitt National Design Museum in New York [20], historical newspapers [9] or literary fiction [13]. Beyond text documents, multimedia objects (e.g.

photos, images, videos) contained in the digitized archives also lend themselves to software-assisted analysis; this meets a growing interest in aiding the described kinds of inquiries in different domains in the humanities [37].

In our work, we have focused on a specific role of photographs as a source of knowledge aiding historical analysis of social ties and relationships. While photographs alone do not suffice to tell us about the exact nature of real-world social ties, they can be used to study the self-representation of historical entities to the public (e.g. politicians, state authorities, governments). On a micro level, image composition, gestures or objects can give insights into their self-understanding and the way in which they wanted the public to understand them. On a macro level, we can ask who was photographed often with whom, which constellations between actors, nations or positions occurred more often than others and how they changed over time. It needs to be stressed that the value of any such approach depends on the quality of the data source (e.g. image quality, available metadata), the homogeneity of the photos (e.g. group photos vs. specific scenes during the event) and so on. Beyond the analysis of people, places and contexts, historians may also use the photos to give a more vivid impression of the events or the interactions between specific politicians. Constructing and visualizing networks of such relationships for aiding research in historical disciplines has confirmed its usefulness for serendipitous discovery of patterns, which can hardly be discovered otherwise [6][12]. Such network visualizations can also serve as powerful tools to illustrate historical processes to non-expert audiences (e.g. changing personal constellations in the process of the European integration). At the same time, while there are a number of approaches to constructing such networks from texts, there is little work on trying to elicit them directly from multimedia objects, such as photos.

III. RELATED WORK

A. *Social network analysis and visualization in historical network research*

Social network analysis (SNA) is concerned with systematically recording and analyzing social structures, based on mapping interactions in groups of individuals into graphs representing individuals (nodes) and the connections between them (edges). This is typically performed to gain insight into the structure of social relationships and their influence on different aspects of the group’s operation, (e.g. actions, decisions, behaviour) or its impact on social contexts (e.g. organizations). SNA has been applied in a number of fields: from social and political sciences, to bioinformatics, computer science, organizational or business studies (see [34] for an overview). While early work involved manual mapping of interactions and relationships (e.g. through interviews, manual document analysis), current approaches analyze digital data, such as Email, and document exchanges [24][32] or interactions in online social networks [19].

Historical sources and databases have always been of interest to practitioners of social network analysis, most of whom applied their expertise in social science methodology

to answer questions also related to the social sciences [27]. More recently, historians have begun to adapt SNA tools and methods to answer historical questions. Their approach differs in that they use formal methods in conjunction with their training in source criticism, leading to a more nuanced dialogue between data, visualization and hermeneutics with the ultimate goal to link their findings back to a historical narrative on specific historical events. Research in the field uses databases and visualizations to capture complex relations between entities such as individuals, organizations, places and objects. Previous works for example covered Early Modern nepotism and marriage strategies, interlocking directorates, credit markets, trade, covert networks and migration [40].

Most current applications in digital humanities apply such techniques to collections of textual documents to extract the relevant entities and relationships between them. Examples include the analysis of letter correspondence, travel and intellectual network data to trace intellectual exchange between scholars in the early modern period [43][38] or the analysis of the correspondence with an individual cultural historian [36]. Some approaches analyze connections between Wikipedia entries to extract and visualize social networks of historical personalities in a specific scholarly area, e.g. identifying central figures in the philosophy domain and their connections within the network [2]. Others have used online genealogical databases to discover and visualize historical social networks of family histories based on document transcriptions [21]. Visualization of social networks has also been applied to the domain of prosopography [28], where text corpora are analysed to “identify the unique individuals who populated a documented social milieu”.

B. Visualization tools and collaborative analysis

Performing such analyses can be aided by general purpose network visualization tools such as Gephi [4], Pajek [3] or NodeXL [30]. They enable interactive exploration of social networks by employing force-directed placement algorithms [16] to generate visual layouts where strongly connected nodes are placed closer together and less-connected ones further apart. This enables explorative discovery of insights, such as the relationships between individuals (e.g. ego-networks, cliques), their roles or effect on group actions (e.g. brokers, bridges, dominant personalities). This is further supported with various measures of structural graph properties (e.g. centrality, in-betweenness, clustering). Such tools provide different levels of detail and interactions based on general-purpose principles of information visualization for managing the switching between graph overview and inspection of individual portions or specific details (e.g. the “overview first, zoom and filter, details on demand” model [29]). They are conceived as information visualization tools and their designs reflect the underlying premise of organizing and displaying complex information structures with the goal of “amplifying cognition of human users” [5]. At the same time - their great contribution to complex analysis of social networks notwithstanding - they do not explicitly address

collaborative scenarios, nor the specifics of the inquiry and knowledge construction processes in the humanities (i.e. in historical research in our case). This contrasts them to recent approaches that seek to align the development of visual tools with the humanist inquiry process, such as the VennMaker [12] and the KNOT tool [32] aimed at using visualization to support *manual* construction of graphs in performing a research inquiry (inspecting documents, discovering relations etc.). It also contrasts them to recent approaches from HCI aiming at coupling visualization with collaborative analysis [18]. Such approaches are closer to the notion of knowledge visualization by supporting the knowledge exchange and communication between users, a point distinguishing knowledge from information visualization [14].

C. Analysing multimedia collections

Similarly, the construction and visual exploration of social networks from multimedia collections has been relatively little addressed in digital humanities. One line of approaches analyzes textual metadata of multimedia elements (photos, images) to construct networks based on co-occurrence relationships [5][8]. Direct analysis of multimedia content (e.g. patterns in colour distribution, textures or image clustering by visual similarity) has been used foremost in cultural analytics [25], archaeology and cultural heritage applications (see overview in [33]). The exploration of multimedia archives has also been reviewed in relation to the specific requirements of humanities scholars towards automatic feature extraction and the linking of image regions and person identities [1] [21]. The idea of combining of human and machine computation has been reviewed for specific use cases, such as the analysis of rock n' roll archives [31]. The use of such methods to support network analysis in historical research however, hasn't yet been directly explored. At the same time, related techniques e.g. for constructing social networks based on analysis of co-occurrences in photographs by means of face recognition have been explored in multimedia information retrieval and HCI (e.g. photo albums [7] or security investigations [15]).

IV. APPROACH AND DESIGN PRINCIPLES

Accordingly, the main goal of our approach has been to develop an interactive tool closely connecting the network analysis and visualization from multimedia collections with the research process of historians. Rather than using a general purpose, off-the-shelf tool, this requires a tailored solution embedded into existing working practices *and* able to facilitate novel approaches. The tool should thus 1) support historians in more effectively investigating familiar kinds of research questions within their research tradition, 2) enable them to ask and investigate new kinds of questions (e.g. not possible with traditional methods) and 3) stimulate new practices in their research (e.g. collaborative analysis of historical multimedia collections).

A. User-centered and interdisciplinary design

Such challenges can only be tackled in a multidisciplinary framework; in our case provided by the collaborative European project CUBRIK [39].

CUBRIK develops novel techniques for multimedia search and analysis that combine human and machine intelligence. To inform the development of such technologies and to demonstrate their practical application potential, the technical research is organized around selected reference domains – one being the Digital Humanities i.e. historical network research. To this end, we have developed an iterative process in which user-centred requirements analysis is performed in a way that aligns the existing needs of history researchers (user-pull) with the identification of opportunities created by novel technologies (technology-push) [35].

This process involved and required active participation of historians, computer scientists, HCI researchers and interface designers (broadly comparable to the approach used in the development of the KNOT tool, recently reported in [32]). In contrast to “traditional” user-centred design practice of identifying requirements and then deriving technology support based on the perceived requirements fit, we gave the same importance to the perspective of the “users” (identifying existing practices and their needs) *and* to the perspective of “technologists” (what novel technology developments could be imagined) – and tried to create a space of shared understanding between them, before pinpointing the “needs” and the new “means”.

In doing so, we combined several techniques in iterative cycles between the two stakeholder groups. Contextual interviews with history experts from the CVCE and a digital humanities expert were performed to understand current research practice, the most important challenges in their work and the potential for computer support. User stories were formulated to translate user needs and opportunities into more tangible ideas, to stimulate the history experts’ thinking about technological opportunities and to reflect on their work from this perspective. Mock-ups were developed to translate and visualize user stories into possible functionalities and tools.

Multiple focus groups and workshops were organized, in which first user stories and later mock-ups were shown and discussed with: 1) users, in order to discover whether the ideas appealed to their (latent) needs, and 2) computer scientists, in order to elicit ideas for technological support, for feasibility, salience and innovativeness from a technology research point of view. The next sections discuss the insights gained, the identified needs and potentials for computer support and the resulting design principles guiding the development of the histoGraph.

B. Understanding historical research: multi-perspective “detective work”

the so-called triad of the historical method: the identification of suitable sources (“heuristics”), their evaluation (“source criticism”) and their interpretation in context to each other (“interpretation”). This is a highly non-linear process. Source identification is done in an iterative manner, with the expert using his own knowledge, asking archivists or colleagues or sometimes turning to web search. When working with photographs, once a relevant photo is identified, experts often search for additional related

materials, such as images, videos, articles etc. For instance: historical articles, an image or an agenda with handwritten notes for a specific meeting. Source evaluation is crucial, as the experts are very critical of the sources: their history and inherent biases must be known (provenance). The interpretation of and coupling of information taken from sources may be characterized as “detective work”. Historians are used to weighing contradicting information against each other and to acknowledging and reconciling multi-perspective (often differing) views on events. They make use of archival materials to formulate or support a hypothesis.

Much of their work is about interpretation, pattern recognition, and identifying causal relations (“connecting the dots”): for instance, when interpreting a series of meetings between people that ultimately led to an important historical event. The work therefore involves sensitivity, context knowledge and intuition. Knowledge – in terms of research findings such as interpretations of the importance, role or relationships of a specific actor or event – is considered as a construction highly dependent on the sources of artefacts explored and the disciplinary framing from which it originates. Data and information itself are not considered as given but already as (subjective) constructs determined by the methods of their collection and analysis – an inherently interpretative stance that is considered at odds with the positivist logic often assumed in the construction of technical systems (for an overview, see [23]).

C. Identifying people, objects and relationships The word “data” is plural, not singular.

For the scope of this paper, two aspects of the interpretation process are of particular interest. First, identifying people, objects or context appeared to be an important yet time-consuming part of their work. This is now done by asking colleagues or peers, by looking for clues in books, or by using an image matching service such as TinEye [46]. In addition, our analysis revealed that historians are already using existing social media networks, e.g. Twitter, to distribute mostly image-related queries such as “Who is this person?” among colleagues. Second, experts expressed a need to visualize the relations between documents in an interactive network, where relations can be added, deleted, or otherwise manually influenced (in line with the needs of digital humanities experts recently reported in [32]).

Finally, an important caveat needs to be stated. Given their training in source criticism and the careful processes of selection and contextualisation, historians are wary of any black box system that remains incomprehensible to them. It is impossible for them to trust a system to make the ‘right’ assessments important to them, in particular with respect to source evaluation: transparency of data processing and user control are key.

Accordingly, the origin of any visualized relationships must be presented in a comprehensible manner, e.g. by linking clearly to content from which they stem. Providing a graph-based overview of the relationships between nodes (e.g. persons) is a good starting point for the inherently explorative process of achieving an increased understanding

of the available materials in the archive, in the light of the their research questions. But rather than merely presenting system-generated graphs, history experts need to be able to interact with a graph-based overview where nodes and relations can be added, deleted, edited or enriched with own information. They need to be able to change perspective by (content-/time-based) filtering and to switch between macro- and micro-scale views of the visualization at will: grasping the overview of the archive as a whole and gaining insight into specific relationships between graph nodes. They need access to related sources that place specific content from the graph (like an image of an event) in a historical and cultural context – and need to be able to trace back the origin of every single piece of information.

D. Obtaining and interpreting information and knowledge from others

The information that is shown as well as the way it has been processed, annotated, ordered, and presented should be explained to the history experts in detail. They should be able to record and share different interpretations, annotations, and explanations related to the data and to export the results of their investigations outside of the visualization tool (e.g. images, source references). The tool should support the historians in dealing with a high level of uncertainty about and incompleteness of the underlying data, e.g. by allowing and presenting different interpretations of the identity of unknown actors or a place where a photo has been taken, by enabling an overview of multiple co-existing versions of a document (e.g. targeted at different audiences). It should be possible to easily submit research inquiries to colleagues or peers: for example, to receive suggestions for the identities of people, buildings or objects in a photograph. It should be possible to view a list of experts that have consulted or annotated a particular document or image, to assess or inquire about relevance of a piece of information. Related information for a selected node (e.g. person, photo) should be provided in a ranked list with a clear explanation of the relevance ranking.

E. Synthesizing design principles

Table 1 presents the synthesis of these lessons learnt into key design considerations that integrate and compare our findings with those from [32], where a comparable methodological approach has been pursued in the realization of a visual tool (see Chapter 3.2), for a similar application domain and for a similar target user group as histoGraph. Even though the two tools fundamentally differ in the specific design and functionalities, the table shows that the *design principles* for both KNOT and histoGraph exhibit a very similar understanding of humanities scholars’ needs. This can be interpreted as a form of (implicit) cross-validation of the results obtained in both cases. Thereby, our results show more room for automated analyses and (contextualized) suggestions in building and visualizing the graph (in contrast to the manual approach of KNOT) and in supporting collaborative analysis and sharing of knowledge.

TABLE I. DESIGN PRINCIPLES FOR NETWORK VISUALIZATION TOOLS FOR DIG. HUMANITIES (INCLUDING RESULTS FROM [32])

Design principle	Our result	Uboldi et al.
Visual interface for interpretative activities	X	X
Contextualizing information/knowledge	X	X
Supporting discovery and (re-) presentation of multiple perspectives	X	implicit
Allowing manual network construction	implicit	X
Understanding and revealing data sources	X	X
Data manipulation and enrichment	X	X
Embedding in existing practice	X	implicit
Transparency of data processing, re-ordering, and presentation	X	X
Support active construction of knowledge	X	X
Support collaborative analysis and sharing of knowledge between scholars	X	

V. THE HISTOGRAPH PROTOTYPE

The developed prototype of the histoGraph tool implements these requirements and design principles into a visual tool for explorative investigation and collaborative analysis of social networks constructed from historical multimedia collections. The prototype is based on the collection of historical photographs (3000 photos) representing the main events and actors in the history of the European integration of the Centre Virtuel de la Connaissance sur l’Europe (CVCE). The main element is an interactive network visualization enabling the exploration and collaborative analysis of relations between persons occurring in the photo collection, accompanied with media objects (text articles, images, videos) depicting their contexts (e.g. events, places).

A. Constructing the social graph

To construct the social graph, the digital collection is processed by an automated workflow (a CUBRIK pipeline [35]) that creates the incidence matrix of a weighted graph representing the estimated social links between the identities found in the collection. The process can be schematized as a sequence of six tasks, executed by a combination of automated processing and human analysis in a coordinated manner (Fig. 1). First, the digital collection is ingested into the system; images are normalized, indexed, and associated with labels extracted from the textual metadata accompanying an image (if any). After that the images are subjected to face detection using an off-the-shelf component [41].

The results can contain both false positives (wrongly identified image parts not representing faces) and false negatives (undetected faces). The detected faces are validated by a crowd using a commercial crowdsourcing market [44]. Since the complexity of the task is low (“Is this a human face? Yes/No”), a general crowd of non-expert users can fulfil it.

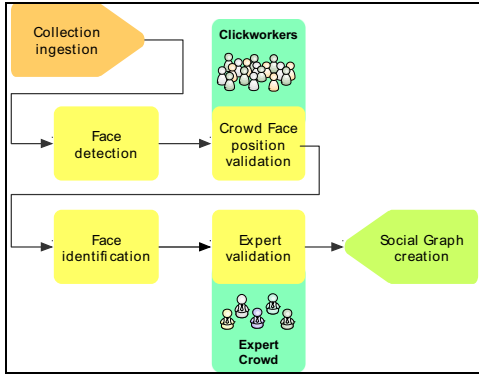


Figure 1. Combination of human and machine computation in the creation of the social graph

To obtain identities of the persons represented by the faces, automated face recognition is performed (using an off-the-shelf face recognition component [41]). As the collection already contains a set of validated portraits of identified historical personalities, these are used for the initial recognition (matching faces with unknown identities to known portraits). This produces a ranked list of potential identities for each detected face in a photograph. The result is a set of faces associated with possible identities of the persons represented in a given photo (with a probability score). Subsequently, the computed identities can be validated by a crowd of experts. To this end, the system allows historians to resolve ambiguous identifications by themselves, or by forwarding a validation task to their social circle, which includes more experts (see 6.5). Accordingly, for each person a co-occurrence measure is computed as the frequency with which person A (all faces identified as person A) and person B (all faces identified as person B) appear in the same photo. Based on such co-occurrence information a social graph is constructed, representing a network of persons (nodes) and connections between them (edges).

B. Visualizing the social network

The resulting social graph is visualized in the form of an interactive network of persons and connections between them, where the strength of the connections (edge width) represents the frequency of co-occurrences of two persons in the photos in the collection (Fig. 2). The visualization is implemented by means of a force-directed graph layout that places frequently co-occurring nodes closer together, while nodes that co-occur less often are placed farther apart (performed by javascript components Crossfilter and D3.js).

The nodes in the visualization represent historical persons, while the graph's edges represent the photographs that connect them. Based on the date, location and event of the photographs (where available), the graph is positioned in its temporal and spatial context. One way for scholars to explore the co-occurrences and possible social relations of the different historical persons is to look at the complex structure of the histoGraph from a birds-eye-view. This view can give new insights into seemingly well-known knowledge (recapitulation) as well as lead to surprises, e.g. identifying unknown persons in photographs by finding unexpected links. Main actors can be identified, paths between actors as well as overlaps between clusters or sub-networks discovered. Users can hover over nodes and see the names of persons and their connections to others highlighted, or filter the graph by time (see 5.3).

C. Exploring and discovering relationships

Clicking on a single node zooms in to the ego-network of the corresponding person (e.g. the ego-network of Margaret Thatcher in Fig. 3). A search bar helps locating specific persons in the graph and the auto-complete suggestions can lead to discovering persons that may remain unnoticed in the exploration. Users can also zoom in and click on the edge connecting two persons to view all photographs from which the co-occurrences were derived and visualized (Fig. 3).

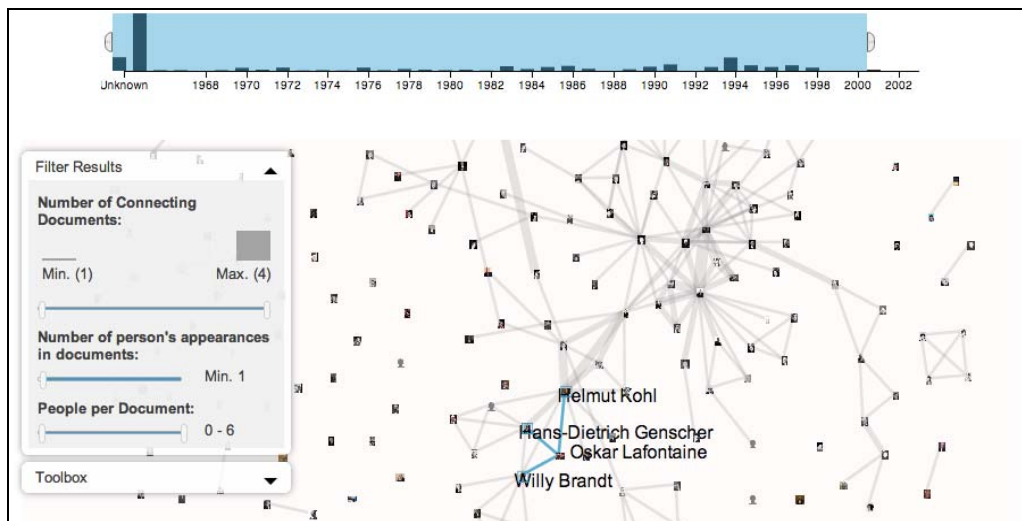


Figure 2. Social graph visualization in histoGraph

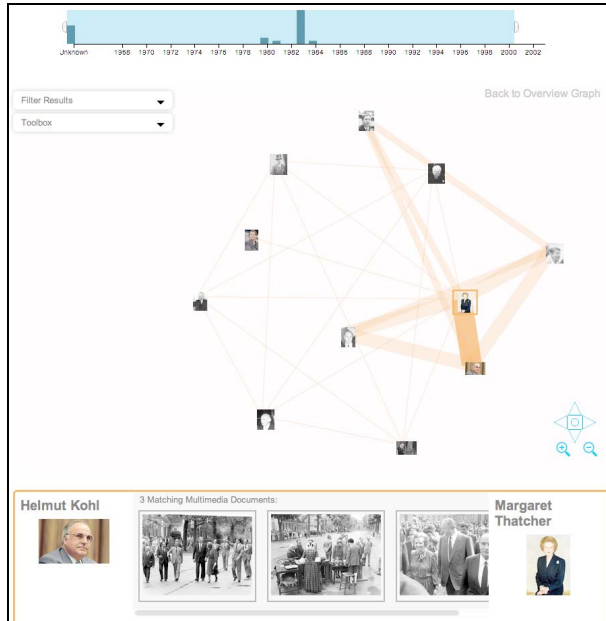


Figure 3. Ego-network zoom-in for in-depth contextual exploration of relationships

This supports the required transparency and provenance of sources and methods used in the visualization. A list of additional documents related to persons and photographs that connect them is also retrieved to reveal more about the context of the relationship and entice further research, e.g. finding out more about a discovered connection between two historical personalities.

Contextual filters such as the timeline and source-based filters enable users to generate smaller sub-graphs based on different parameters. The timeline allows the user to specify a time period that should be considered for building the social graph - and visualizes how many images are available as a basis for this operation in a year-by-year histogram (Fig. 2, 3). This is another means for assessing the relevance of the displayed results. The source-based filters include e.g. the “number of connecting documents”, that can be used to specify the minimum and maximum number of co-occurrences between two persons that must exist for a link between them to be established. This allows for filtering of weak ties (e.g. showing only connections made up by one or few co-occurrences) or strong ties (showing only connections made up by “many” co-occurrences). Weak ties can help identify new surprising connections, while strong ties allow for the identification of the main actors. Similarly, the “people per document” filter can be used to constrain the analysis to photos of small or large groups (e.g. to focus on bi-/trilateral or on multi-lateral meetings only). In a similar manner, place-based and event-based filters can be provided, when such data becomes available (e.g. through crowdsourcing and collaborative inquiries described in the next sections).

D. Verifying and creating new relationships

The visualized relations are initially based on co-occurrences that are calculated using automated face detection and recognition algorithms. Some of the depicted persons are manually confirmed based on available metadata from the CVCE collection, but still many depicted persons are not correctly identified yet.

With an annotation tool embedded into the histoGraph, the identities of persons can be verified or corrected manually by the scholars through a kind of implicit expert-based crowdsourcing. This is important because historical photo collections contain photos of different quality (e.g. size, resolution), shot from different viewpoints and angles and from different time periods - which makes the automatic face recognition difficult and unreliable [10]. Moreover, while important personalities can be more easily identified due to likely availability of named reference portraits for the automatic recognition, identifying less well-known persons and their social context may be just as important for the researchers. Accordingly, an annotation tool allows users to correct the automatically determined identities or to suggest new ones for unidentified faces (Fig. 4).

Visualizations of humanities data are often man-made and therefore prone to be incomplete and ambiguous [11][17]. In writing, scholars acknowledge these shortcomings by a nuanced use of language. Data visualizations however tend to be suggestively absolute: what they show appears to be true without a doubt, what they do not show appears to be non-existent.

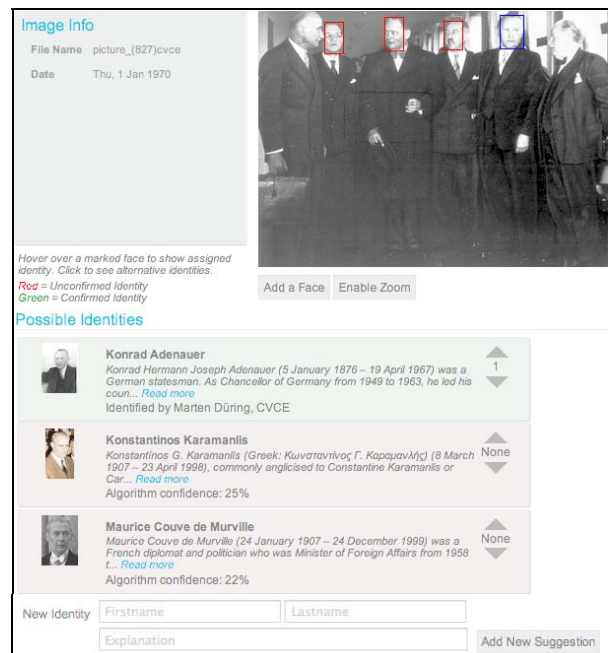


Figure 4. Correcting and suggesting identities through expert-based crowdsourcing

A critical use of such visualizations therefore requires background knowledge of the creation of data, their provenance and the technical implementation. To ensure transparency and provenance in our case, annotations such as new identity suggestions are accompanied by an explanation from their authors, so that other users can reconstruct their reasoning or look up sources (Fig. 4). The tool also incorporates a voting mechanism similar to established practices of question and answer sites (e.g. Stackoverflow [45]). We use this to enable discussions about uncertain or contradictory information, common in the humanities. Rather than limiting annotations to a struggle over one contested truth, users can actively reflect discussions by voting identity suggestions up or down, depending on how likely they find them. After an identity is verified or corrected, the co-occurrences matrix is recalculated and the graph is adjusted to reflect the new findings.

In addition to correcting and verifying identities of depicted persons, it is also possible to expand the graph further by creating new nodes (persons) and relationships (co-occurrences). With the annotation tool, person faces not yet detected in photos can be manually marked by the users and subsequently identified (through automatic face recognition and manually by the scholars). When a new person is identified, a new node and related co-occurrence relations are added to the graph. Users can also add new photos (or photo collections) and annotate them. The graph accordingly expands, enriched with additional nodes and relations with each added photo.

E. Inquiring and sharing expert knowledge

To facilitate a more explicit knowledge exchange between humanities scholars, users of the histoGraph can initiate explicit research inquiries that are distributed to their peers (the expert-crowd) within the application, via E-Mail or Twitter. A research inquiry links to a specific photograph and can refer to the identity of a depicted person, contextual information and other metadata about the photograph, such as: “Who is the marked person?”, “Where was this photo taken?”, “Was this the event x?”, “What role did this person play here?”. Replies by other scholars feed back into the application and are directly visible to all users, in order to support the sharing and construction of knowledge through collective action [26]. The new pieces of knowledge accumulated in this way extend the social graph: they can result in new nodes and relationships and enrich the available contextual data (e.g. places, events) that expands the range of possible exploration, filtering and search operations.

F. Gaining insights, discovering hypotheses

An example of the kind of insights that historians can gain through the described visual exploration and interaction is given in Fig. 5. A first analysis of the histoGraph network by a historian reveals that individuals tend to cluster together by regional vicinity. Figure 5 shows an annotated screenshot displaying how politicians and diplomats from the Netherlands and Belgium form a cluster; similarly Italian and Spanish representatives tend to co-occur.

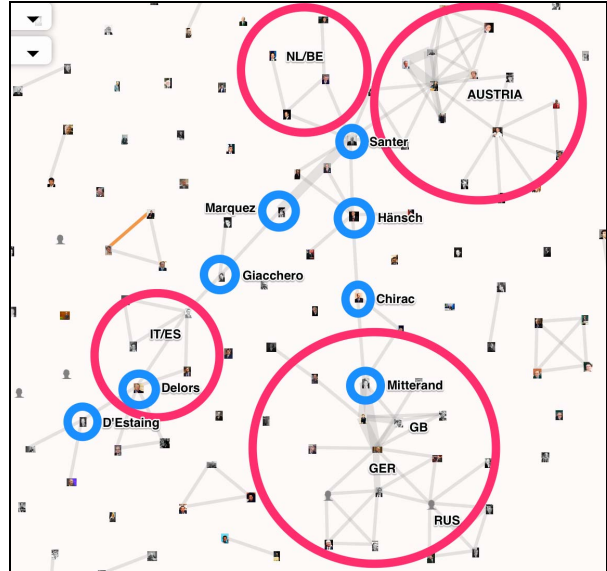


Figure 5. Insight: regional (or cultural) clusters?

Germany is closely connected to the UK (mainly through Kohl and Thatcher) and Russian representatives only occur in relationship with German chancellor Willy Brandt. French statesmen Jacques Chirac and Francois Mitterrand take middle positions between the different clusters and the high-ranking EU representatives like Klaus Hänsch. We need to be careful not to draw premature conclusions from these observations and need to revisit them once additional photos have been added to the graph. But even at this stage, this network representation raises questions: How can we explain that there is no French cluster? Why are there no ties between Germany and Austria?

VI. PRELIMINARY EVALUATION

A preliminary evaluation of the current histoGraph prototype was conducted with humanities scholars working in the field of European integration at CVCE (5 participants). The participants were given a short introduction to the prototype and completed three scenario-based tasks related to their research practice. Feedback was collected in one-on-one debriefs after each task, via observation, through a questionnaire and a group discussion at the end. Overall, the participants found it useful for their research to be able to explore photo-based connections between historical persons with the histoGraph, acknowledging the tool as an efficient alternative to their current work processes for discovering such relations (4 out of 5 agreed or fully agreed to both statements, Fig. 7). They stated that they would use the tool to explore the evolution of relations over time and based on other contextual impact points like key historic events; to explore the impact of such relations in larger networks, to identify persons as well as to study networks, co-occurrences, connections, affiliations and allegiances. One participant remarked that he could discern groupings roughly corresponding to basic political leaning (“left” vs. “right”).

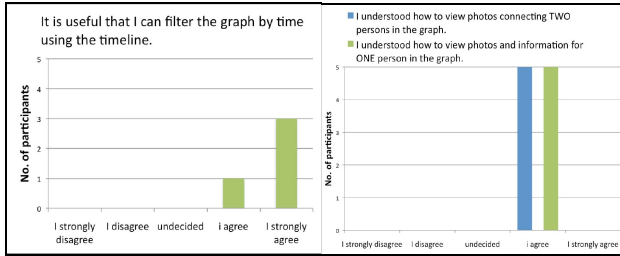


Figure 6. Evaluation: crucial functionalities

The participants also found the structure of the graph logical and understood how to explore nodes and edges between two persons to view the photos connecting them (Fig. 6). Including contextual information was identified as a crucial requirement with special emphasis on the temporal filter (“the timeline was most important”).

The participants stated that being able to influence and make changes to the graph was important for the usefulness of the tool (“need ability to correct or mark false information in the graph“, “irritates when a relation is missing that you know should be there”), confirming the requirement that users should be able to add their own relations and nodes to the graph (Chapter 4). This is already supported in the current prototype (nodes are added when new persons are identified in photographs, and edges are added and corrected as more persons are identified correctly by the users). However, such feedback indicates that this implicit approach is not yet clearly recognized as a means to influence the graph and should be highlighted more explicitly in the interface. This could even be used as a way to motivate users in the future to add their own document findings to the existing collections and to annotate photographs, as the effect of their actions on the quality and usefulness of the graph is made clearly visible.

Some usability issues were identified in the evaluation, including the somewhat ambiguous naming of the source-based filters, the size of the displayed thumbs for the nodes (too small) and the need for more cohesiveness between the histoGraph visualization and its annotation tool. The issue that was perceived by participants as the most critical was the (insufficient) relevance and quality of related text documents that are retrieved for each person and connection (in addition to the photographs that construct the actual

graph). While text documents as such were appreciated as important source for further researching possible contexts of connections, their relevance, ranking and sources were criticized as unclear. This again underlines the importance of provenance in this context, requiring transparency of sources, querying and ranking rules.

VII. CONCLUSIONS

The histoGraph prototype and its design process show how the construction and visualization of historical social networks from multimedia collections can be accomplished in a way that supports historians in the discovery and historical analysis of relationships between people, places and events. The described interdisciplinary collaboration allowed us to identify specific classes of requirements and design principles for building tools that embed into existing practices of humanist inquiry and can stimulate new approaches.

First evaluation results suggest that the applied combination of automatic analysis, network visualization and crowdsourcing provides useful support for the explorative investigation of social relations from historical photo collections. Qualitative user feedback also suggests that the proposed design manages to bridge the gap between established network visualization techniques and traditional hermeneutic research methods in historical research. It highlights the crucial importance of system transparency and provenance information and points to critical issues such as the need for transparent query and ranking mechanisms of accompanying documents revealing more possible contexts for the graph connections. Next steps involve the integration of specific network analysis functions (e.g. centrality, in-betweenness), the extension of the graph with place/ event detection and filtering as well as the extension of co-occurrence analysis to other media types (e.g. Wikipedia entries, text documents, news videos). The audio-visual archive of the European Commission will be integrated to extend the collection and a more detailed evaluation in form of a pilot study is planned.

VIII. ACKNOWLEDGEMENTS

This work has been partially funded within the EU FP7 project CUBRIK (grant agreement n° 287704).

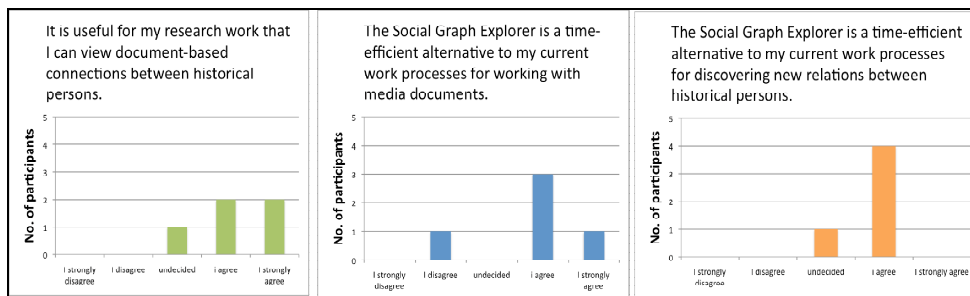


Figure 7. Evaluation: task-adequacy and improvement of existing work practice

IX. REFERENCES

- [1] R.B.N. Aly, et al., "Link Anchors in Images: Is there Truth?" Proc. of 12th Dutch Belgian Information Retrieval Workshop (DIR 2012), Ghent, BE, Feb 2012.
- [2] S. J. Athenikos, and X. Lin., "WikiPhiloSofia: extraction and visualization of facts, relations, and networks concerning philosophers using Wikipedia," Conf. Abstracts of Digital Humanities 2009, 2009, pp. 56-62.
- [3] V. Batagelj, and A. Mrvar, "Pajek-program for large network analysis," Connections, vol. 21(2), 1998, pp. 47-57.
- [4] M. Bastian, et al., "Gephi: an open source software for exploring and manipulating networks," ICWSM, 2009, pp. 361-362.
- [5] B. Berry, "Friends for Better or for Worse: Interracial Friendship in the United States as Seen Through Wedding Party Photos," Demography, vol. 43(3), 2006.
- [6] T. Brughmans, "Networks of Networks: A Citation Network Analysis of the Adoption, Use, and Adaptation of Formal Network Techniques in Archaeology," Literary and Linguistic Computing, vol. 28(4), August 2013, pp. 538-562.
- [7] M. Crampes, et al., "Visualizing social photos on a hasse diagram for eliciting relations and indexing new photos," Visualization and Computer Graphics, IEEE Transactions on, vol. 15(6), 2009, pp. 985-992.
- [8] D. J. Crandall, et al., "Inferring Social Ties from Geographic Coincidences," Proc. National Academy of Sciences of the USA, vol. 107(52), Dec 2010, pp. 22436-22441.
- [9] G. Crane, and A. Jones, "The challenge of virginia banks: an evaluation of named entity analysis in a 19th-century newspaper collection," Proc. of the 6th ACM/IEEE-CS joint conf. on Digital libraries, New York, NY, ACM, 2006, pp. 31-40.
- [10] M. Dionisio, et al., "Building social graphs from images through expert-based crowdsourcing," Proc. of SoHuman 2013, pp. 29-34.
- [11] J. Drucker, "Humanities approaches to graphical display," Digital Humanities Quarterly, vol. 5(1), 2011.
- [12] M. Düring, et al., "VennMaker for Historians: Sources, Social Networks and Software," Redes. Revista Hispana Para El Análisis de Redes Sociales, vol. 21(8), 2011.
- [13] D. K. Elson, et al., "Extracting social networks from literary fiction," Proc. 48th Annual Meeting of Assoc. for Comp. Linguistics. Stroudsburg, 2010, pp. 138-147.
- [14] M. J. Eppler, and R. Aslak Burkhard, Knowledge Visualization, 2011, pp. 987-999.
- [15] M. J. Eppler, and Roland Pfister, "Best of Both Worlds: Hybrid Knowledge Visualization in Police Crime Fighting and Military Operations," Proc. of ACM CIKM 2013, ACM, 2013.
- [16] T. MJ Fruchterman, and E. M. Reingold, "Graph Drawing by Force-Directed Placement," Software: Practice and Experience, vol. 21(11), 1991, pp. 1129-1164.
- [17] R. Häußling, Visualisierung sozialer Netzwerke, VS Verlag für Sozialwissenschaften, 2014.
- [18] J. Heer, and M. Agrawala, "Design considerations for collaborative visual analytics," Information visualization, vol. 7(1), 2008, pp. 49-62.
- [19] J. Heer, and D. Boyd, "Vizster: Visualizing online social networks," Proc. of IEEE Symposium on Information Visualization, IEEE, 2005, pp. 32-39.
- [20] S. van Hooland, et al., "Exploring Entity Recognition and Disambiguation for Cultural Heritage Collections," in Literary and Linguistic Computing, 2013.
- [21] F. de Jong, et al., "Audio-visual collections and the user needs of scholars in the humanities: a case for co-development," Proc. 2nd Conf. on Supporting Digital Humanities (SDH 2011), Centre for Language Technology, Copenhagen, Denmark, 2011, p. 7.
- [22] D. J. Kennard, et al., "Linking the past: discovering historical social networks from documents and linking to a genealogical database," Proc. of the 2011 Workshop on Historical Document Imaging and Processing, ACM, 2011, pp. 43-50.
- [23] C. Lorenz, Konstruktion der Vergangenheit: eine Einführung in die Geschichtstheorie, Böhlau, 1997.
- [24] D. MacLean, et al., "Groups without tears: mining social topologies from email," Proc. 16th international conf. on intelligent user interfaces, ACM, 2011, pp. 83-92.
- [25] L. Manovich, "Cultural analytics: Visualizing cultural patterns in the era of more media," Domus, vol. 923, 2009.
- [26] J. Novak, Human Computation and Collective Action, in Pietro Michelucci (ed), Handbook of Human Computation, Springer, Nov. 2013.
- [27] J. F. Padgett, and C. K. Ansell, "Robust Action and the Rise of the Medici," American journal of sociology, 1993, pp. 1259-1319.
- [28] P. Schmitz, and L. Pearce, "Berkeley prosopography services: ancient families, modern tools," Proc. 1st Intl. Workshop on Collaborative Annotations in Shared Environment: metadata, vocabularies and techniques in the Digital Humanities, ACM, 2013.
- [29] B. Shneiderman, "The eyes have it: A task by data type taxonomy for information visualizations," Proc. IEEE Symposium on Visual Languages, IEEE, 1996, pp. 336-343.
- [30] M. A. Smith, et al., "Analyzing (social media) networks with NodeXL," Proc. of the 4th international conf. on communities and technologies, ACM, 2009, pp. 255-264.
- [31] C. G. M. Snoek, et al., "Crowdsourcing rock n'roll multimedia retrieval," Proc. of the international conf. on Multimedia, ACM, Oct 2010, pp. 1535-1538.
- [32] G. Uboldi, et al. Knot, "An interface for the study of social networks in the humanities," Proc. of the Biannual Conf. of the Italian Chapter of SIGCHI, ACM, 2013.
- [33] C. Warwick, et al., Digital humanities in practice, Facet Publishing, 2012.
- [34] S. Wasserman, Social network analysis: Methods and applications, vol. 8, Cambridge Univ. Press, 1994.
- [35] L. Wieneke, et al., "Building the social graph of the History of European Integration. A pipeline for Humanist-Machine Interaction in the Digital Humanities," 1st Intl. Workshop on Histoinformatics, Springer, 2014.
- [36] Burckhardt Source url: <http://www.burckhardtsource.org/>
- [37] Bomb Sight. Mapping the WW2 census url: <http://www.bombsight.org/#15/51.5050/-0.0900>
- [38] ePistolarium url: <http://ckcc.huygens.knaw.nl/>
- [39] EU FP7 CUBRIK url: <http://www.cubrikproject.eu/>
- [40] Historical Network Research url: <http://historicalnetworkresearch.org/>
- [41] Kee Square url: <http://www.keesquare.com/>
- [42] League of Nations url: <http://www.lonsea.de/>
- [43] MRoL url: <http://republicofletters.stanford.edu/>
- [44] Microtask Oy url: <http://www.microtask.com/>
- [45] Stackoverflow url: <http://stackoverflow.com/>
- [46] Tineye url: <http://www.tineye.com>