Visual Informatics 5 (2021) 45–60

Contents lists available at ScienceDirect

Visual Informatics

journal homepage: www.elsevier.com/locate/visinf

Gone full circle: A radial approach to visualize event-based networks in digital humanities

Velitchko Filipov [a,*], Victor Schetinger [a], Kathrin Raminger [b], Nathalie Soursos [b], Susana Zapke [b], Silvia Miksch [a]

[a] TU Wien, Favoritenstrasse 9-11, 1040 Vienna, Austria
[b] Musik und Kunst Privatuniversität, Johannesgasse 4A, 1010 Vienna, Austria

ARTICLE INFO

Article history:
Received 24 November 2020
Received in revised form 15 January 2021
Accepted 19 January 2021
Available online 1 February 2021

Keywords:
Event-based networks
Information visualization
Digital humanities

ABSTRACT

In the application domain of digital humanities network visualization is increasingly being used to conduct research as the main interests of the domain experts lie in exploring and analyzing relationships between entities and their changes over time. Visualizing the dynamics and different perspectives of such data is a non-trivial task but it enables researchers to explore connections between disparate entities and investigate historical narratives that emerge. In this paper we present Circular, an interactive exploration environment to visualize event-based networks and support research in digital humanities through visualization of historical subjects in space and time. Our radial design is the result of iterative collaboration with domain experts, and we discuss the process of collaborative development and exploration of public music festivities in Vienna as an example of immersive development methodology. We validate our approach by means of both domain and visualization expert interviews and show the potential of this approach in supporting the visual exploration of historical subjects. We discuss our design rationales, visual encodings, and interactions as to allow the reproducibility of this approach within a framework of transdisciplinary collaboration with digital humanities.

1. Introduction

In recent years, academic and applied research in Visual Analytics (VA) has seen remarkable progress that helps advance the field and grow the community. VA has been increasingly applied to a multitude of different fields and disciplines and established itself as a direct interface between the data, the users, and their tasks (Miksch and Aigner, 2014). VA leverages the potential of human cognition to make large amounts of data and information tractable. While digital tools and modern computational methodologies are increasingly being used across most fields and research disciplines, in the humanities it is still quite common for most research to be conducted in archives gathering material sources. This is mostly due to the fact that research questions and methodologies from the humanities stem from a different epistemological standing compared to traditional technical fields (Hinrichs et al., 2017). New technologies are changing the landscape of this domain by making research materials widely accessible and enable scholars to explore their data, find patterns, and present their work. A big challenge in this domain is modeling data for historical research properly and accurately. As stated by Börner et al. (2019) many artifacts and phenomena in the humanities are difficult to model as variables, due to the abstractions necessary to make these commensurable.

The data in the domain of digital humanities can be seen as rich interpretable data (Lamqaddam et al., 2018) exhibiting relational, spatial, and temporal characteristics, evolving and changing over time. The general idea is that the cultural and social phenomena that can be observed in such disciplines, based on the data available, does not occur naturally, but is in fact shaped by different influences, interests, and power relations of certain actors that change over time. The intricacies and complexity of this type of data can be represented by a network (Schetinger et al., 2019). Networks are one of the most flexible structures to model such phenomena and can illuminate dimensions and scales of historical events that are otherwise difficult to conceptualize (Lincoln, 2016). The core analysis tasks supported by network visualization that are of main interest for the application domain of the digital humanities are to identify and understand relationships between actors, to find and highlight central or similar objects, to observe the underlying structure of the network, to reveal outliers or deviations from the expected network behavior, and to observe the changes occurring to the network over time.

* Corresponding author.
E-mail address: velitchko.filipov@tuwien.ac.at (V. Filipov).

https://doi.org/10.1016/j.visinf.2021.01.001

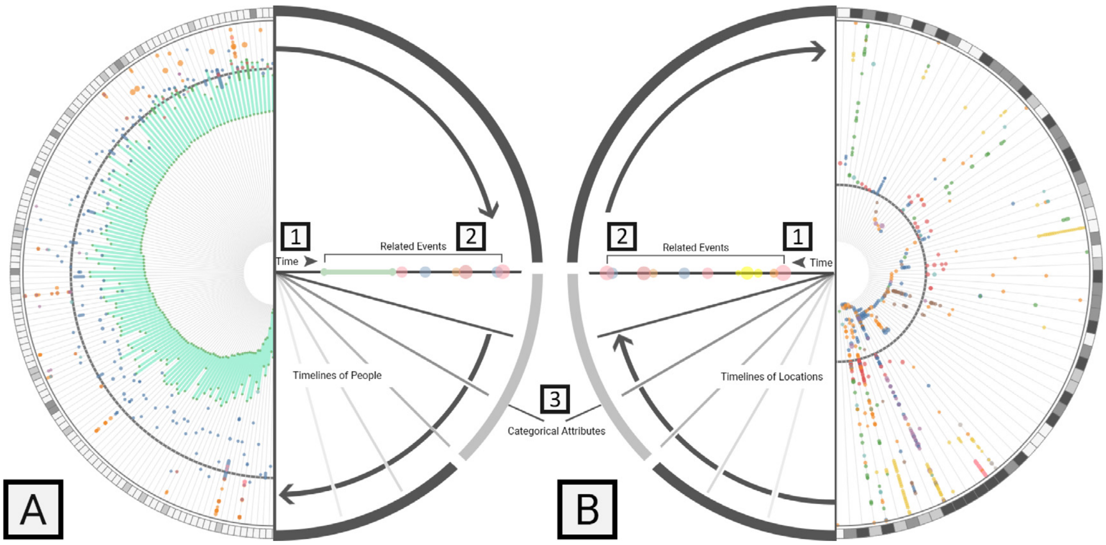

Fig. 1. Circular depicting two networks from our data (A and B). The persistent entities and their temporal developments are encoded as rays along the circle (1), whereas their related events are represented by dots superimposed on the respective rays and are color-coded according to their association to different themes (2). On the exterior of the circle we encode categorical attributes that can be user selected (3). (A) depicts people, their related events, and the themes associated with those events. The rays are ordered according to user selected criteria, in this case by date of birth. The people of interest in this dataset have a large amount of events related to them occurring after their death. (B) shows important locations in Vienna alongside events that have been organized there and the themes associated with those events. In this dataset we can observe the change of popularity for certain locations over time and their shift in motifs.

Network visualization has proven to be a successful tool to model data and phenomena for such tasks, yet most modern solutions do not account for the dynamic nature of the data. Real-world data is rarely static and for most applications and problems it is of essence to model and visualize the evolution or change in the network, its actors, their relationships, and movement over time. The dynamics of such a network can be expressed by event-based network visualization and we consider the challenge of visualizing this type of data for historical research.

This paper presents the results of interdisciplinary collaboration in the project Interactive Music Mapping Vienna (IMMV[1]). The project aims at mapping the history of public festivities in Vienna during the Second Republic (from 1945 to present day), exposing narratives that contributed to the construction of the city's musical identity and the connection to historical events. It stands in the domain of musicology, a field in which aesthetics play an important role in motivating research questions, and is conducted mainly by historians, who prime for critical analysis of events and their representations. As VA researchers our challenge was to provide efficient and effective means for the humanists to explore their subjects with the desired level of rigor and artistic inquiry.

We present a circular visualization that is flexible and transparent to the underlying subject. It uses event-based networks as inputs, considering the existence of a small set of entities in space and time: people, places, events, and themes. This ontologically minimal model is powerful enough to represent many facets of historical subjects and explore trends and narratives mainly through their temporal and spatial developments, serving as VA support to humanists. Our solution was developed through an iterative process of immersive collaboration (Hall et al., 2020) that was shaped by the continuous research of material sources and construction of the data. The evaluation was designed by the humanities domain experts with the methodological aid of VA researchers. Our main contributions in this paper are:

- Characterizing the problem domain of musicology and musical history, along with the data and tasks the domain experts have for conducting their research.
- Exploring the design space of radial visualizations, the conceptualization, and prototypical development of Circular - a radial visualization of event-based networks.
- Elaborating the evaluation results of our user study along with feedback and insights that were gained with Circular from both domain and visualization experts.

2. Problem domain characterization

The main focus of the interdisciplinary research project IMMV is how music acts as a social identification instrument in the urban context and how music is functionalized to urban symbolic politics. The concrete subject of research is selected festivities in the public space of Vienna during the Second Republic, starting with 1945 up to the present day.

Our approach, Circular, is influenced by our initial attempt to analyze how the City of Vienna honored its exiled musicians (Filipov et al., 2019). This was an exploratory, open-ended task on uninvestigated data which required visualizing multiple timelines, exploring relationships between people and events, and outlining pattern formation. The core idea was to leverage adjacency between timelines to invoke gestalt effects (Wertheimer, 1923). The same timelines can be reordered and regrouped according to different criteria, e.g., birth dates and honoring events, providing representations that highlight meaningful aspects for domain experts. However, this first design was a first sketch of the overall vision and restricted to exiled musicians. In this contribution we abstracted the conceptual design to be domain independent, applicable to various objects (entities) in time and space, such as locations. We oriented our design according to the data, users, and tasks (Miksch and Aigner, 2014) as well as the nested model by Munzner (2009). Consequently, this abstraction allows to utilize Circular to visually explore and analyze similar problems of different research fields, which obey to the same problem characteristics.

2.1. Data

Our data represents a network relating disparate entities and embedding them in space and time. The data is dynamic and is

[1] http://www.musicmapping.at/ (accessed Oct., 8, 2020).

shaped as new entries are added or new sources are discovered and analyzed. The main difficulties stem from the fact that most information on the subject is not easily accessible or even available as digital assets. Therefore, the construction of the data itself can be considered original research and visualizing the coverage of information to orient research efforts is a crucial task. At the time of writing, our database comprises of a large amount of entities that are all related among each other through events. We model these entities as nodes in a network with the relationships between them constituting our links. Our data consists of six different types of entities (and their respective counts): Events (1243), Historic Events (78), People/Organizations (1538), Themes (61), Locations (180), and Sources (1279) with a total of 15,866 links between them. In our use case and considering our application domain networks are a suitable data abstraction to model historical research and can showcase the topological structure of these highly interconnected entities, with events being the most central nodes in the network tying together different types of nodes (people, organizations, themes, locations, and sources) and embedding these in spatio-temporal frames of reference. In this specific paper, the data being depicted is only a subgraph of our entire data, representing a single theme ("Musicalization of the Topology" see Fig. 2-B).

Initially, our data was based on a standard model for describing digital resources and offered little flexibility to our domain experts for modeling and conducting their research. Furthermore, the initial model did not offer support for one of our main tasks, which was relating different entities. Throughout the design and development we iterated on this data model, implementing relationships and adapting entities and their properties as needed (see Fig. 2-A for the final model). This process concluded in six main types of entities, acting as nodes (each with their own set of specific properties and attributes), that can be related to each other through individual events, that tie these together in spatial and temporal frames of reference. An example of this could be a musical festival that was held at a specific location, organized by a group of people with different political affiliations, and sponsored by a specific organization. This festival could further be related to a multitude media sources, such as photographs, video or audio material, news paper snippets, etc.

The data model offers significant advantages over our domain experts' previous method for modeling historical research, specifically our approach offers flexibility to model the complex intricacies of historical narratives (i.e., by relate prominent entities to these) and enforces scope. To understand why one must realize that the main task the domain experts have is analyzing material sources and structuring these to expose historical narratives and to present the complexity of their historical research. From a single material source many entities can be derived, e.g., from a news paper article one can compile a set of people, an event, a location, and a date. On the other hand, from an encyclopedia one can gather a more exhaustive list of entities and their relationships. Prior to our approach the data model did not allow to establish and construct relationships between existing entities and our domain experts ended up modeling an excessive amount of information from the sources that does not explicitly enforce the desired subject: the role of music in an urban context as an instrument of societal identification and the question of how music is to be functionalized for constructing the image of the City of Vienna. Furthermore, our approach offers flexibility to our domain experts and enables them to easily interface with and curate the data. The data is directly tied to the visualization and can be considered an online approach, in that it allows for real-time updates. This provides better support to our domain experts in analyzing, exploring, and identifying how the overall shape and scope of the data changes between subsequent modifications, such as addition, deletion, or updates.

By using events as our central ontological objects, carefully sampling key points throughout Vienna's history, we are able to create spatio-temporal frames of reference and visualize historical contexts between people, places, organizations, and themes. A more comprehensive and detailed explanation of the data and data model can be found in our previous work (Filipov et al., 2019).

Our data can be described as an event-based network. Event-based network data consists of a set of events over a period of time, each of which can relate to multiple objects (O'Madadhain et al., 2005; Simonetto et al., 2020). Traditional dynamic networks in most cases aggregate the network to a single snapshot or multiple timeslices, often resulting in a loss of information as opposed to event-based networks that model events with real time coordinates (Du et al., 2017; Monroe et al., 2013; Monroe et al., 2013).

2.2. Users

Our approach was designed with expert users in mind, i.e., humanists with domain knowledge that intend to perform in-depth historical research on specific narratives and topics. Domain experts in the field of digital humanities mainly conduct their research by gathering material sources from archives, analyzing collections of textual material, and structuring all this information in a comprehensive manner to establish narratives, present their work, and expose their ideas. In their research methodologies and activities VA solutions are not common, this presents an opportunity for novel VA solutions in this domain. We intend to support them by creating VA approaches for analyzing and visualizing the data, exploring narratives or possible trends, investigating relationships between disparate entities, and conveying spatio-temporal patterns that might emerge. In our evaluation we had two groups of users participate: domain experts, with specific knowledge centered around the data and time period of interest and visualization experts, with experience in information visualization methods and VA techniques.

2.3. Tasks

To understand the problem domain in more detail and best provide support for our domain experts, their tasks, and current challenges, there was an extensive period of discussion about the construction of the data model, the data and scope, and their research questions. The experts suggested a lot of domain related literature as a way to introduce us to their ongoing research and the challenges they face. Furthermore, this became a way for us to become involved in their domain, gain a better understanding of their data and requirements, and to explore new domain-inspired visualization design spaces that presented interesting research challenges from a VA perspective as well but also have a contribution in the problem domain. These were necessary steps in order to immerse ourselves in the problem domain and familiarize ourselves with their research (Hall et al., 2020).

Our domain experts were interested in achieving a variety of tasks ranging from high-level to low-level as outlined in the multi-level typology of abstract visualization tasks from Brehmer and Munzner (Brehmer and Munzner, 2013). Starting from high-level tasks, an emphasis was put on presenting results, discovering insights, and producing narratives from the data. Some examples of these are questions posed by our domain experts that refer to how the motifs and themes of locations where events are held have changed over the course of time and how the City of Vienna treated and honored exiled musicians that returned to their homeland. Mid-level tasks included supporting the experts

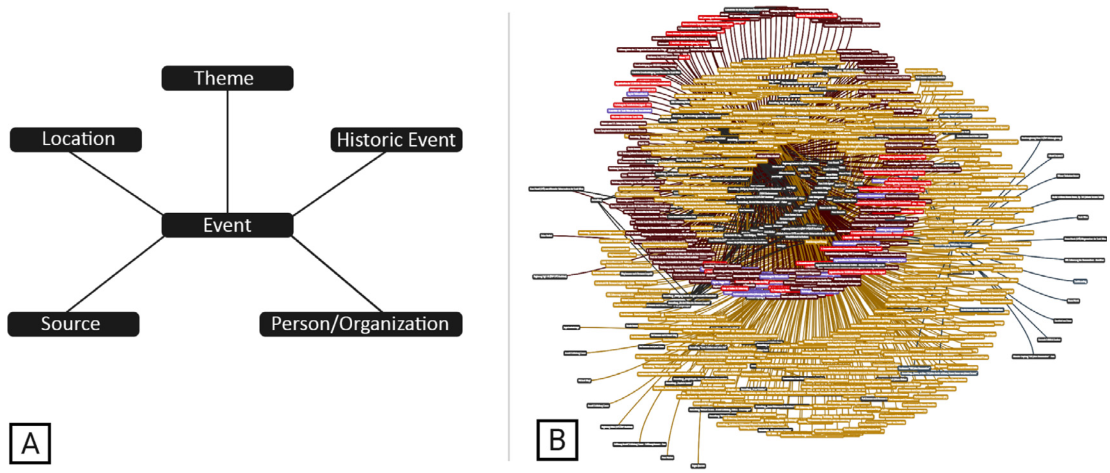

Fig. 2. An overview of how our event-based network is modeled with real examples. Events are the most central object that can relate people, organizations, locations, sources, themes, and historic events. In (A) we present the topology of our network and in (B) the subgraph that represents the data related to a specific theme ("Musicalization of the Topology") that is used throughout this paper depicted as a node-link diagram.

in free-form browsing of the data, as well as having the functionality to search for specific entities or browse for elements of interest in the visualization. Our experts were interested in being able to lookup specific people, events, and locations to find which were more important, influential, or well connected, as well as having the ability to explore and find elements of the visualization that were pointing towards interesting periods of time, such as events related to many other events or people or people and locations with irregular event patterns throughout their lifetime. Finally, for the low-level tasks the intent was to enable our experts to query elements of the visualization for more details on their specific characteristics or relationships with other entities. The questions our domain experts have that fall into this category include, identifying specific events, people, or locations, discover details about them and their relationships to other entities in the dataset (the full set of questions can be seen in Table 1). These tasks and data of our experts became the basis for our domain problem characterization.

There has been significant work on establishing task taxonomies for visualization and graph design and analysis (Tory and Moller, 2004; Andrienko and Andrienko, 2006; Lee et al., 2006; Schulz et al., 2013; Munzner, 2014; Ahn et al., 2014), however, we found the multi-level task typology by Brehmer and Munzner (2013) to be better suited for our specific problem domain. The authors outline that the typology allows for complex tasks to be expressed as a sequence of interdependent simpler tasks and this notion provides us with a flexible methodology for conducting our research and developing our approach. Following the nested model approach by Munzner (2009) we abstract these specific problems, questions, and data from their domain into a more abstract and generic description which resulted in an illustrative list of tasks in accordance with the multi-level task typology by Brehmer and Munzner (2013):

T1: Present - Our approach should support a straightforward, aesthetic, and engaging interface to present and gain an overview of the event-based network. Additionally, it should enable our domain experts to communicate the results of their research by telling a story with the data. The core goal of this task is for the users to work with the visualization in order to communicate information, tell a story with the data, or guide the audience through the data and visual representation. Essentially, our approach should enable users to convey the general shape of the data, but also provide different interactions to slice, reshape, reorder, and view different modes of the network. The presentation and dissemination of historical research is one of the main objectives of the domain experts. Examples of questions from our domain experts that fall into this category include, e.g., *"In which locations was the First of May event celebrated in 1924?"* or *"Can you name 3 exiled musicians that do not have a street in the city named after them?"*.

T2: Compare - Our domain experts expressed their interests in comparing entities or groups of entities that exhibit similar characteristics to observe their development over time. The intent of this task category is to support them in discovering different groupings of the network based on specific properties of the entities, such as, e.g., grouping locations by district or people by their profession and supporting them in observing their similarities or differences. Some questions that our experts specifically had are, e.g., *"Which location was used most often to celebrate the First of May between 1918 and 2018, and which one was used the least?"*, or *"Are there more male or female people in the current network?"*.

T3: Explore - In addition to the presentation of research results and interactions that enable the experts to slice and compare parts of the network, our domain experts found it important to have interactions that enabled them to navigate and explore the network in more detail. This task category is concerned with the ability to search for data points with specific characteristics and investigating their properties and relationships in depth. This task starts from an overview level of the visualization, exploring a large set of data points (i.e., people or events) and includes, e.g., searching for outliers and anomalous or periodic patterns in time. This task is supported by methods for navigating, zooming, and panning. Examples of questions that motivated this task category are, e.g., *"Which events took place at the Befreiungsdenkmal? Which music was played there? Who organized the events? Which historic events are they referring to?"* or *"Can you find the location(s) attributed to Regenbogenparade in 2016?"*.

T4: Identify - As our domain experts are concerned with manually entering the data, shaping, and scoping it, being able to identify continuities and discontinuities of themes throughout the historical developments of the City of Vienna becomes an interesting topic. The experts are very detail-oriented and one of their main research challenges is to identify and summarize population level trends and patterns that arise and decline over time. Such patterns are inherent to the data and become apparent at different granularities. Identifying these patterns and trends become interesting directions for further research in the problem domain (i.e., groups of people based on their gender, profession, or people that have been exiled versus those that have not). This task differs from *T3: Explore*, in that the set of targets or data

Table 1
A table of questions along with their corresponding task category. The questions have varying degrees of difficulty that the participants performed throughout the user study. It is important to note that questions can appear in more than one task category.

Question	Task Categories
Q1: At which locations was the First of May celebrated in 1924?	T1
Q2: When did the first event related to "Stadtbranding" happen?	T3
Q3: Which themes are related to the events taking place at "Stadthalle"?	T1, T2, T3
Q4: When is Erich Wolfgang Korngold born / When did Erich Wolfgang Korngold die?	T5
Q5: Which events are related to Alban Berg?	T3, T5
Q6: Can you find an interesting event OR person in the timeline OR circle? Inspect it in more detail.	T1, T3
Q7: Which location was used most often for "Politische Inszenierungen" during the Second Republic?	T2, T4
Q8: Which events took place at the "Befreiungsdenkmal"? Which music was played there? Who organized the events? Which historic events are they referring to?	T1, T3, T5
Q9: Which location was most used to celebrate the First of May between 1918 and 2018, and which one was used the least?	T2, T3, T4
Q10: Can you find the location(s) attributed to "Regenbogenparade" 2016?	T3, T5
Q11: In which year(s) did the opening of the "Wiener Festwochen" NOT take place at "Rathausplatz"? Which location(s) were used instead?	T2, T3
Q12: In which district(s) did fewest events take place?	T1, T2, T4
Q13: Which exiled musicians never came back? Name two.	T2, T3, T5
Q14: Can you name 3 exiled Musicians that do not have a street named after them?	T1, T3, T5
Q15: Are there more male or female people in the visualization?	T1, T2
Q16: Can you rank the groups by role based on the number of people in each group.	T2
Q17: Can you name the musician from each group that was born first?.	T1, T2, T4
Q18: Can you find the two conferences? When did they take place and which musicians were related to them?	T2, T3, T5

points that have been found or discovered through exploration can be further investigated and the user can drill-down to their properties and relationships. In this case the goal is to identify, for specific data points, important characteristics and Andrienko and Andrienko (2006). Some questions that fall into this task category are, e.g., *"Which location was used most often for political stagings during the Second Republic?"* or *"In which district(s) did the least amount of events take place?"*.

T5: Verify - Finally, a task category that is not present in the multi-level typology by Brehmer and Munzner (2013), but proved to be essential for the domain experts was the ability to verify and validate the data. It is crucial to ensure that the data being depicted is correct, as all data is entered manually and any erroneous data points, entries, or irregularities could point to entries that need to be corrected or provide interesting clues to open questions for historical research. Questions that belong to this task category are, e.g., *"When was Erich Wolfgang Korngold born/When did Erich Wolfgang Korngold die?"*, *"Which events are related to Alban Berg?"*, or *"Which exiled musicians never came back?"*.

3. Related work

3.1. Radial visualization

Visual representations of data that are based on circular shapes are referred to as radial visualizations (Burch and Weiskopf, 2014). Historically they date back to the nineteenth-century, with a most famous example being the Rose charts introduced by Florence Nightingale (Nightingale, 1987) and the introduction of the pie chart in 1801 by Playfair (1801). Circular designs, however, have been used consistently throughout human history across different cultures, from mandalas to navigational charts, generally associated to cyclic and periodic phenomena. Recently radial representations have seen a return in popularity (Shi et al., 2018; Maçãs and Machado, 2018; Castermans et al., 2019), presumably due to their natural shapes, aesthetic look, and memorability (Borkin et al., 2013). Draper et al. (2009) provide a comprehensive survey on radial representations and present taxonomy based on seven design patterns. In their work the authors investigate problem domains where radial visualizations are applied and typical usage domains include: hierarchical structures, relationships among disparate entities, ranking of search results, and serial periodic data. Accordingly, some examples of radial representations applied to these domains and different datatypes, include set-typed data (Alsallakh et al., 2013), multivariate data (Bale et al., 2007), hierarchical data (Burch and Diehl, 2008), and time-oriented data (Keim et al., 2004). An interesting and widely adopted application of radial representations in the domain of genomics has been explored by Krzywinski et al. (2009) where the authors present their approach, Circos, a technique utilizing a circular ideogram layout for visualizing and comparing the complexities of genome relationships.

Evaluations of radial visualizations have focused mainly on comparing radial and linear representations of periodic time-oriented data (Waldner et al., 2019), assessing the strengths and weaknesses of radial visualizations in comparison to Cartesian representations (Diehl et al., 2010), and outlining the benefits and drawbacks, as well as providing guidelines for when radial visualizations are more suitable for given tasks and datasets (Burch and Weiskopf, 2014). Despite all the drawbacks of radial visualizations, there are also strengths in visually presenting data in a radial fashion and this is confirmed in the work of Borkin et al. (2013), where they describe radial visualizations as being more aesthetic, natural, and memorable. This notion is further supported by the work of Hohman et al. (2020), where the authors elaborate that an audience which finds the content presented in an aesthetically pleasing manner is more likely to have a positive attitude towards it, engage with it, and learn from it.

Aesthetics also play an important role for us and our application domain as the main goals include fostering engagement, driving curiosity, exploring, and explaining historical narratives and connections in the data.

3.2. Event-based visualization

Since the main subjects of our data are events, their dynamics, and their relationships, we also look into the related fields of event-based analytics and timeline visualization. Monroe et al. (2013) explore how an entire event-based dataset can be aggregated and transformed so that researchers and analysts can observe population-level patterns and trends. Plaisant et al. (1996) propose a general technique for visualizing summaries of personal histories and other types of biographical data. The authors mention in their work that multiple facets of the data can be reflected in a single overview, using multiple timelines, with icons indicating discrete events and line color and thickness illustrating relationships between events. Cappers and van Wijk (2017) argue that current approaches focus only on either the temporal analysis of a single property or the structural analysis of multivariate properties. The authors propose an approach to interactively define a set of rules using multivariate regular expressions and use glyphs to encode the event sequences. Brehmer et al. (2017) review related work and literature, analyzing over 200 timeline visualizations, and propose a design framework based on three main dimensions: representation, scale, and layout. This work represents the state of the art in timeline visualization and the design framework can be applied to biographical and prosopographical data visualization. Recently, Di Bartolomeo et al. (2020) have evaluated the effects and influence of different timeline shapes (radial and linear) on task performance for temporal event sequences.

3.3. Dynamic network visualization

In digital humanities it is becoming increasingly common for phenomena to be modeled as a network, as this is a flexible structure relating entities with different properties. The different entities can be represented as nodes in a network, for example, actors and events, related among each other and embedded in a spatio-temporal frame of reference. Network analysis can assist researchers in the humanities in exploring dimensions and scales of the data that are otherwise difficult to conceptualize (Lincoln, 2016; Schich et al., 2014). For this purpose, we also present related work from the field of network and graph visualization.

For a comprehensive survey on the state of the art of dynamic network visualization we refer to the work of Beck et al. (2014). In their analysis the authors outline open research challenges and future directions for dynamic network visualization, including visual scalability, hybrid visualizations, and new metaphors or approaches for visualizing the graph and its dynamics. In recent work (Simonetto et al., 2020; Liu et al., 2015) we can see a paradigm shift in the way that time is modeled in networks: from time-slicing and aggregating the temporal dimension to representing it as an event-based network where the nodes and links can have real-valued time coordinates. With time being such an important dimension in dynamic graphs, the problem of identifying patterns and trends in the network becomes of interest. Dang et al. (2016) propose a novel visualization technique called TimeArcs to enable users in exploring patterns in networks over time. With their approach the authors aim to display the evolution of entities as they change over time, highlight related entities by positioning them close to each other, and reduce the amount of line crossings leading to occlusion and visual clutter. Elzen et al. (2014) present, Massive Sequence Views, a technique for the analysis of the temporal and structural facets of the network along with node reordering strategies to find patterns and trends in the data. As outlined in the state of the art (Beck et al., 2014) visual scalability and hybrid approaches present open research challenges, this topic has also been explored by Vehlow et al. (2015) in the context of visualizing group structures in graphs.

In our work we aim to explore new metaphors for representing time in networks. We investigate how we can utilize the design space of radial visualizations and apply them to event-based network visualization. In the following we present our approach, Circular, our design rationales and decisions, the visual encoding of our approach, and tie the features and interactions to the tasks they support.

4. Circular

Circular is our approach for visualizing event-based networks in the application domain of digital humanities. We discuss the design rationales leading up to the final design, the visual encoding that is used throughout the prototype, and the interactions that enable users to interface with the network and its components. Circular is an open-source web application, the implementation and a live demonstration of the approach can be found at the following URL: https://immv-app.cvast.tuwien.ac.at/biographical.

4.1. Design rationales

4.1.1. Designing for musicological research

During the design and conceptualization phase of our approach, we conducted workshops and meetings with our collaborators where we explored visual representations of datasets exhibiting similar properties and discuss the advantages and disadvantages of each approach. We encouraged our domain experts to find and select a few approaches they found aesthetically pleasing, engaging, and informative. The majority of the solutions we discovered were applied to the presentation of historical narratives and biographical data. The approaches were also developed with engagement and learning in mind as they were aesthetically pleasing and contained interactive and engaging visual representations of the data. We proceeded to investigate the visual encodings, layout, and interactivity of each approach and identified the tasks and requirements that were also backed by our domain experts.

An important question we needed to answer was how this data would be visually represented. As we have a large amount of event-based relational data (see Section 2.1), we needed to consider which aspects and properties were important for the tasks our domain experts have. We also consider how we can visually present this information in an aesthetically pleasing, but also functional, and interactive approach. Interactivity and aesthetics were important factors in guiding our design and development, as presenting the content in an aesthetically pleasing manner increases the chances that a user will have a positive attitude towards it, engage with it, and learn from it (Hohman et al., 2020). Considering the density of our network, traditional node-link approaches do not seem suitable here, as there is a significant amount of occlusions, visual clutter, and edge crossings that arise (see Fig. 2-B) and further encoding the temporal dimension in such approaches presents additional difficulties. Therefore, we considered alternative representations and new metaphors for visualizing networks, such as e.g., Mashima et al. (2012), where neighboring regions on the map symbolize related entities in a network. Furthermore, the proposed solution needs to offer a set of features to reshape, slice, filter, and explore the data, based on the interest(s) and task(s) of the user(s) (see Sections 2.2 and 2.3). Our goal was to create an interactive visual prototype that offered users the option of re-configuring the state of the visualization to match their current interests, needs, or tasks, but also presented the data in an engaging and memorable way.

Throughout the workshops with our domain experts we identified that a common theme amongst most visualizations developed for the digital humanities, were circular ideograms, depicting certain datasets and incorporating different ways of encoding the temporal dimension. This notion was further supported by historical examples of radial visualizations that were used to attract a broader audience as their shape and structure are more aesthetically pleasing, natural, and compact compared to their Cartesian counterparts (Nightingale, 1987; Playfair, 1801). Our final design rationale is motivated by the domain experts we collaborate with and presents an interesting opportunity to explore new domain-inspired research and design spaces.

4.1.2. Design considerations

For this purpose, we explore the design space of radial layouts, how they can be applied to event-based network visualization, and consider how we can visually represent the different types of entities in our dataset, their attributes, relationships, and temporal developments. Our main focus was on representing the time-oriented nature of the data, as to enable and showcase pattern and trend formation throughout different points in Viennese history. With these considerations in mind, we investigate

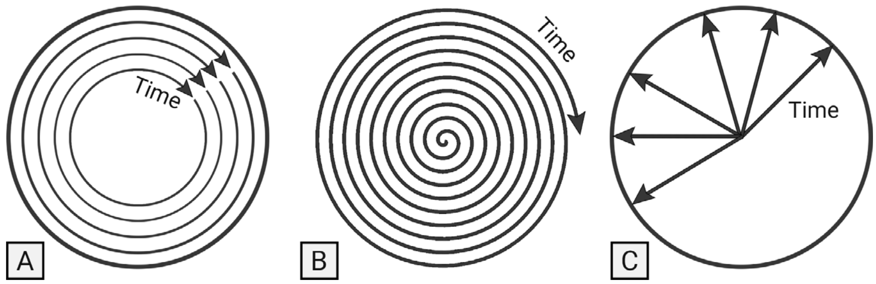

Fig. 3. The three main representations for encoding time in radial visualizations (Draper et al., 2009). (A) Ring-based Time — time is represented as progressing along the circumference; (B) Space-filling — time is represented as a curve starting from the center and growing outwards; (C) Polar Time — time is represented as growing outwards from the center in a ray-like fashion.

the related works of radial visualizations, where there are three primary ways to encode dynamic data in circular layouts (see Fig. 3): (A) Ring-based Time; (B) Space-filling Time; and (C) Polar Time (for a comprehensive survey of the design space of radial visualization techniques, we refer the reader to the work of Draper et al. (2009)). Each of the options has certain strengths and weaknesses, in the following we discuss these and present our preferred solution. We follow the categorization of the radial design space introduced by Draper et al. (2009).

(A) Ring-based Time: This approach encodes time in a circular manner, the evolution of time is presented as going along the circumference from an arbitrary starting point (see Fig. 3-A). There are some drawbacks to this approach compared to the others. First, this layout relies heavily on encoding information using angle and arc length, which can not be perceived very well compared to other perceptual tasks (Cleveland and McGill, 1984). Second, when attempting to visualize multiple tracks, each would need to be represented as a concentric circle, where the elements drawn closer to the circle's center would receive less drawing area compared to those drawn closer to the boundary, skewing the visualization in a way that objects represented closer to the circumference would be perceived as more important. This is true for most radial visualization techniques as the drawing area increases with the radius and elements closer to the center will receive less drawing area, but this is also true for their Cartesian representations as the drawing area is finite.

(B) Space-filling Time: This fashion of encoding data is mostly utilized in visualizing serial periodic time-oriented data. The data in this encoding is visualized as a curve emanating from a point (the center), revolving around its origin as it moves further and further away (see Fig. 3-B). This approach, similar to the ring-based time, also has certain drawbacks, such as a inequality in the drawing areas close to the circle's center and border, but the main issue here would be encoding multiple tracks. We could, for example, have multiple spirals superimposed on each other with each depicting the properties and behavior of a single object, but this has its limitations when it comes to the upper bound on how many objects could be visualized at the same time.

(C) Polar Time: For Polar Time encoding the center usually carries some special meaning and the distance between the center and a node implies some significance. In this approach time is represented as growing outwards from the circle's center to its border (see Fig. 3-C). The temporal information in this approach is encoded along a line, which is an improvement over the other two options that encode time along the circle, since here we are dealing with line length, a perceptual task that can be judged rather well in comparison (Cleveland and McGill, 1984). Furthermore, encoding the temporal facet of the data in this manner offers us freedom to explore what additional dimensions of the data we can encode in other parts of the circle. For example, with this design we can represent multiple objects as neighboring lines and encode further properties on the exterior or interior of the circle (something that was not possible with the other two options). A drawback of this approach is, as previously mentioned, the limited drawing area close to the center of the circle, something we intend to address with interactivity, but we argue that this is still the most balanced approach.

The final encoding we decided on was option (C) Polar Time. We believe this approach is the best suited one for our tasks (see Section 2.3) and the most balanced one, when comparing its' benefits and drawbacks. Our design choices are further backed by the recent results of a user study evaluating the influences of timeline shapes for temporal event sequences (Di Bartolomeo et al., 2020). The authors conclude that linear timelines are preferred by most of the participants in their study and support reading the timelines and temporal sequences of data quicker than non-linear shapes. These claims also support our decision for selecting (C) Polar Time as the temporal sequences of events in this case are depicted in a linear fashion. In the following section we will go into more depth about how we visually represent and embed the data described in Section 2.1.

4.2. Approach

As we mentioned in the data section (see Section 2.1), we have historical data in the form of multiple disparate objects (people/organizations, events, locations, themes, sources, and historical events) that are related among each other and embedded in spatio-temporal frames of reference. We model our data as an event-based network (Simonetto et al., 2020; Liu et al., 2015), where the nodes of the network exhibit spatial and temporal characteristics. Different types of entities are associated with a specific meaning and importance and this should also be taken into consideration when visually encoding and representing the data. An overview of the full prototype can be seen in Fig. 4. In the following we will discuss the visual encodings and interactions in more detail, as well as present some interesting examples from our dataset.

4.3. Visual encoding

When we look at the particular characteristics of people, organizations, and locations, we notice that these are entities that persist through time and define their own life span. These characteristics were taken into account and such objects are encoded as rays along the circle with highlighted features, such as a life spans and important dates. An illustration of how we visualize a single person's or location's life line can be seen in Figs. 5 and 6, respectively. Additionally, related events are superimposed as points on the respective persistent entity's life line, these events have a time span of their own, defined by their start and end dates and additional properties, such as relationships to themes, sources, and historic events. If a single event is related

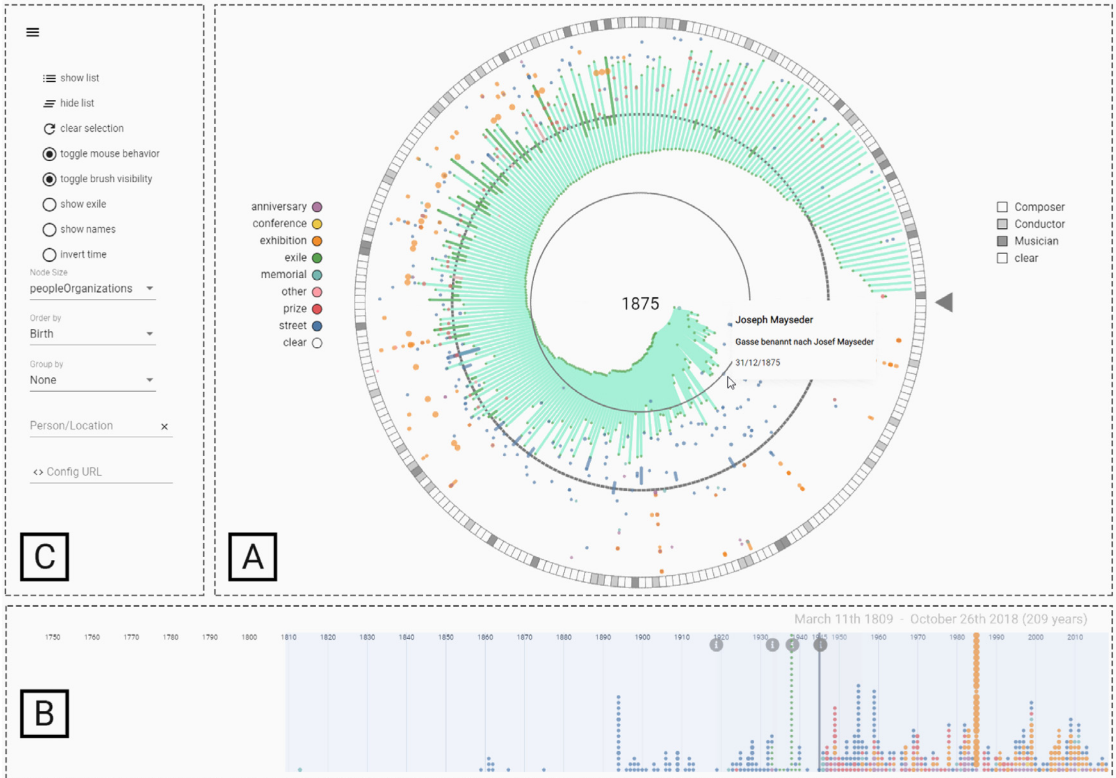

Fig. 4. An overview of our prototype and its components. (A) The main radial view where people, locations, their events, and thematic changes are visualized; (B) The timeline that is used to visualize historic events and ease the users' selection of periods of time; (C) A panel with settings and controls to modify the state of the visualization.

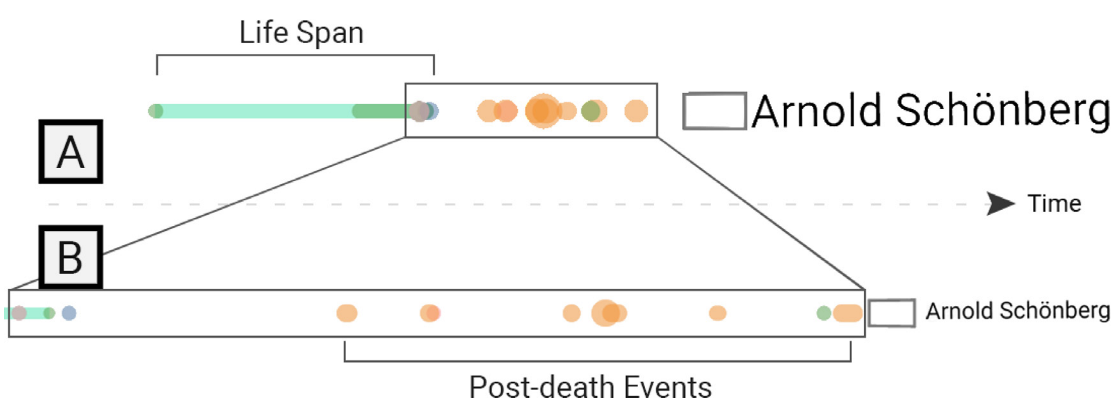

Fig. 5. Arnold Schönberg's lifeline, along with related events and motifs. In this case the events are color-coded according to a classification provided by the domain experts. (A) Here we can see the default temporal granularity and a more scaled up view can be seen in (B). In this example we notice Arnold Schönberg was exiled (darker green bar) close to the end of his life and the bulk of events (orange circles) honoring this person happen after his death.

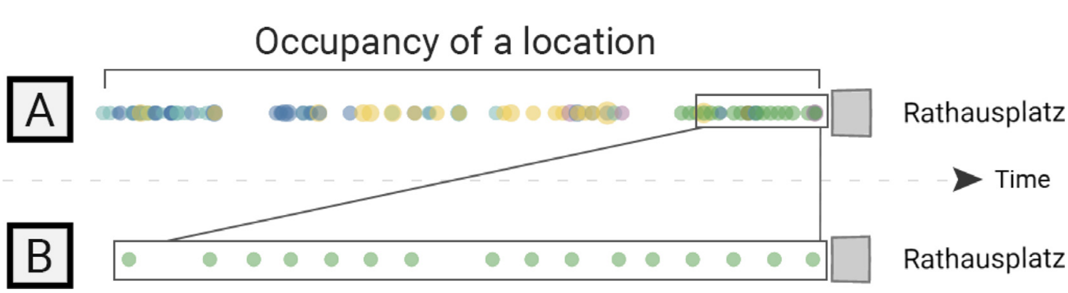

Fig. 6. The lifeline of the Viennese Town Hall Square "Rathausplatz" along with related events and their motifs. In this case the events are color-coded according to the classification of themes provided by the domain experts. (A) We can see the default temporal granularity and a more detailed view can be seen in (B). In this example we can notice how the trends related to events happening at this location is changing over time (shift from blue to yellow to green circles).

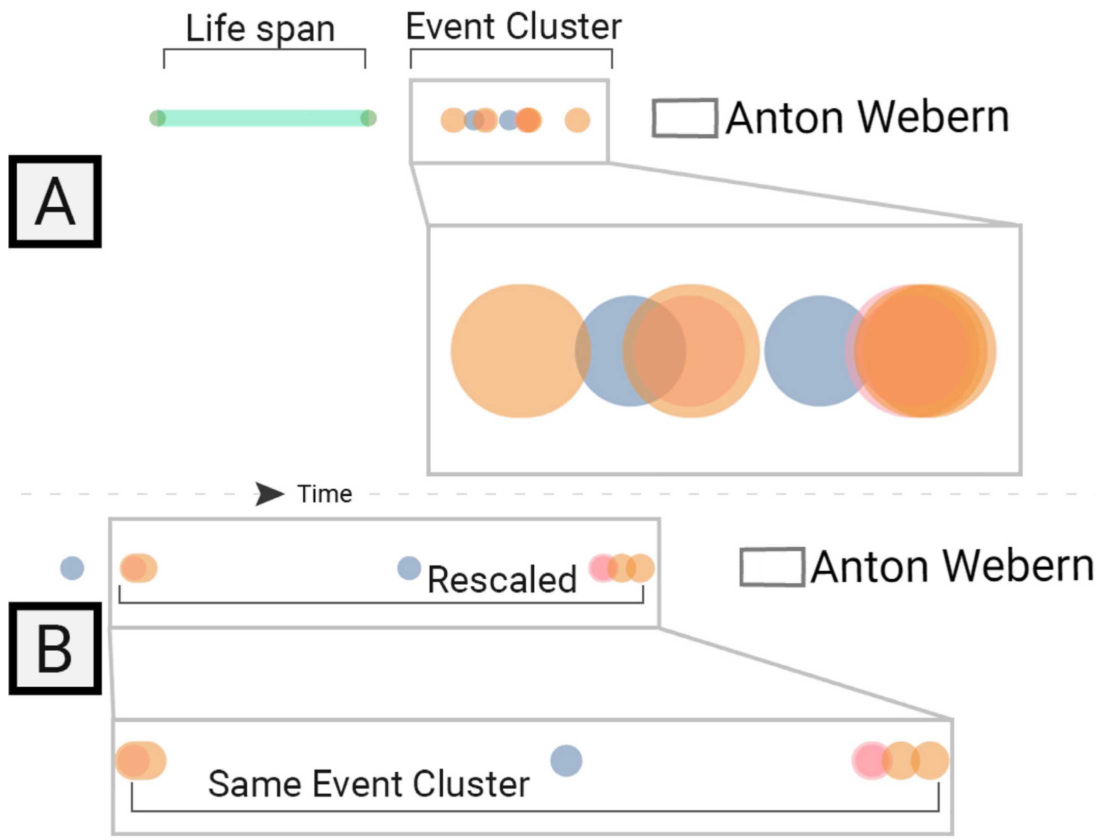

Fig. 7. In (A) we can see an event cluster happening after Anton Webern's death. These events cause clutter and overlaps. In (B) we have scaled the temporal granularity up and the events are better spread out in order to mitigate dense areas and overlaps. In this example the events are color-coded based on a classification provided by our domain experts.

to multiple persistent entities, this event is duplicated along the corresponding rays in the radial component (related to *T1: Present, T4: Identify, and T5: Verify*).

Depending on the frequency of events related to the entities, clusters can be formed. These can be seen as very busy areas superimposed on the person's or location's ray and might result in visual clutter (see Figs. 5-A, 6-A, and 7-A). To tackle this issue, each event is rendered with reduced opacity, which enables the user to see overlapping events, and the size of the points is used to encode the amount of relationships (i.e., node degree) it has to other types of entities such as events, locations, people, organizations, sources, or themes. The nodes' size can be interactively changed based on the user's interest(s) and task(s). We provide a few options to do so, that in turn highlight events of different importance and relevance. This increases readability of clusters and their implied network topology. The node's size guides users towards parts of the network that could be of interest to them based on their selection. For example, if a user is interested in events that are related to a lot of people, they could select the node sizing by person or organization and the changes would immediately be indicated in the visualization. This could imply that those events were considered to be fairly important in Vienna's history and with this selection the size of the nodes would be updated to reflect that information, i.e., events with more related people would be significantly larger than events with less related people (related to *T1: Present and T4: Identify*). Further options to mitigate overlaps and clutter are to allow the user to interactively change the scale or temporal granularity so that the events are better spaced out in time to minimize overlaps (see Figs. 5-B, 6-B, and 7-B).

We make the relationships between events and themes explicit by exploiting the most prominent visual cue — color. Each event is color-coded according to the main theme it is related to, and this makes thematic changes over time visible as seen in Fig. 6, which shows that at earlier points in time this specific location was mostly used to organize events related to the blue theme, then transitioning to yellow, and being dominated by mostly green in present day. This transition of colors reflects how the motifs of events being organized at a specific location change over time. Another example focused on people can be seen in Fig. 5 where after the person's death the bulk of events related to that person are linked to the orange theme. This makes sudden emergence or disruptions of patterns or trends throughout the City of Vienna's historical development visible and provide a potential narrative that encourages the viewer to examine this period of time in more detail (related to *T1: Present, T3: Explore, and T4: Identify*).

Often multiple entities are of interest and a common and important task is to compare sets of entities with common properties to identify similarities, differences, or trends. To encode persistent entities present in the network, we order them as rays along the circle (see Figs. 1-A for people, 1-B for locations). The ordering along the circle is, by default, according to the first event that happens related to each person or location (i.e., birth date). This forms a spiral pattern that implicitly shows the progression of time, as events related to every subsequent person or location happen later (related to *T2: Compare and T4: Identify*).

We experimented with different grouping and ordering strategies and supplemented the visualization with multiple options of combining these that in turn make different patterns stand out (see Fig. 9), each having a different impact on how the user interprets the story told by the data. Along the exterior of the circle, anchored to each ray, are rectangles that are used to encode categorical information about the persistent entities. Such

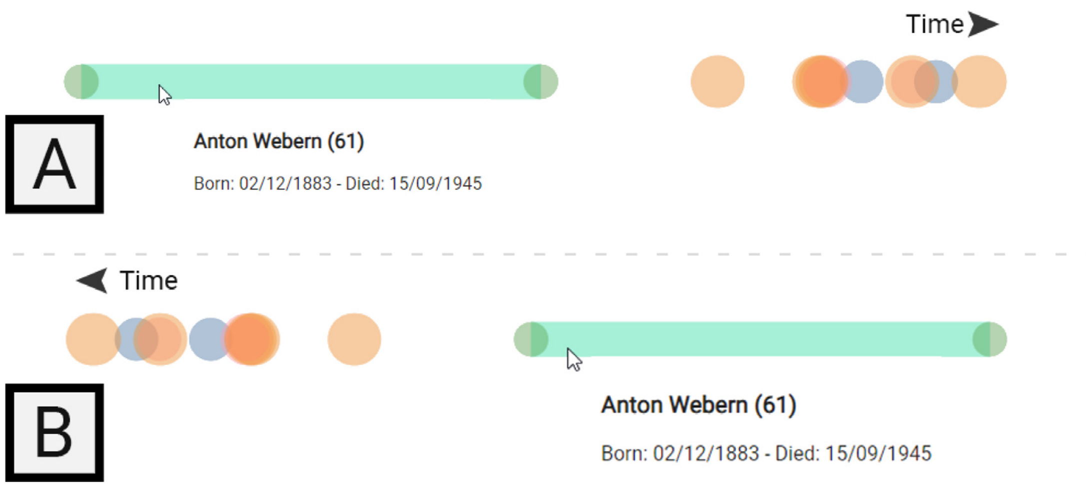

Fig. 8. In Circular we also provide interactions to change the direction of time to resolve overlaps in dense areas close to the center of the circle. In this case the events are color-coded based on a classification provided by our domain experts. In (A) we can see the progression of time is going from the left to the right and in (B) we can see that the direction of time is inverted.

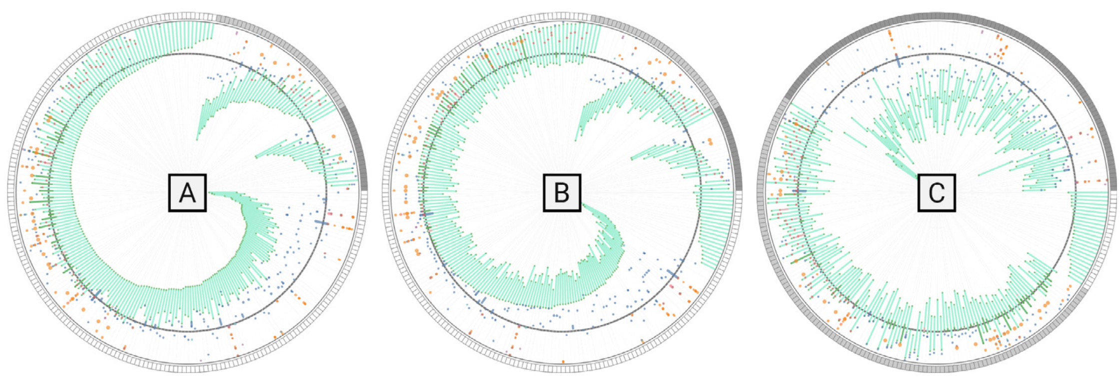

Fig. 9. Three different patterns of people are visualized, each using a different combination of ordering and grouping criteria. (A) People are ordered by their birth date and grouped by their role (composer, conductor, musician). A spiral forms growing outward, which is complementary to that of (B), where people are ordered by their death date and grouped by role and completely different to (C) where people are ordered by their honoring time and grouped by their role to highlight musicians that were honored early on by the City of Vienna.

categories could be, for example, the profession, gender, or period of activity of people, or the type of venue and district of locations. We can use this categorical information to group entities accordingly and outline similarities, differences, or trends within and between these groups. When a grouping option is selected the persistent entities along the circle are arranged in slices alongside others that have the same categorical information. This enables the domain experts to compare different slices of the network that are formed and explore what they have in common or how they differ. The users are free to select properties to group and order people or locations by, thus allowing them to explore the data according to their own interests (related to *T2: Compare, T3: Explore and T4: Identify*).

4.4. Interactions

We offer the user a set of interaction techniques to assist in selection, filtering, ordering, grouping, and inspecting the details of the network. In this section we will present the main components and interactions of Circular.

As depicted in Fig. 4, the central part of our visualization is the radial representation of the event-based network (see Fig. 4-A). In this view we display persistent entities (people, locations) and the events they are related to. The events also reflect relational information as they are color-coded to represent their relationship to different themes and superimposed on timelines reflecting their relationships between different people or locations. The dark gray line that is visible on the interior of the circle denotes the year 1945, which was marked by our domain experts as a baseline in history, used for comparison of the data. To the left side, we have a panel with different settings (see Fig. 4-C) that control the state of the visualization and enable or disable certain interaction techniques. Such controls are, for example, the direction of time (see Fig. 8), ordering and grouping criteria (see Fig. 9), and temporal zooming (see Fig. 7) that we will discuss in more depth.

Our decision for designing and developing a radial visualization for this type of data is mainly motivated by the application domain and our domain experts. We were further interested in exploring the design space of radial visualizations for a few reasons: (i) space is utilized differently compared to rectilinear layouts and offers a more compact embedding of the data; (ii) it presents an interesting opportunity to investigate how we can visualize event-based networks in a radial layout and utilize the inner and outer parts of the circle for more features; (iii) it is important for our application domain to provide an interactive interface that offers an engaging experience in an aesthetically pleasing way and we believe a radial design supports this notion.

As discussed in the design rationales (see Section 4.1), there are certain drawbacks to radial visualizations, with the main concern being the limited drawing area closer to the center. For this purpose, we have looked into different ways of resolving this issue, which lead to a combination of approaches, including:

temporal zoom (see Fig. 7), temporal inversion (see Fig. 8), and panning and zooming. Additionally, we have supplemented the radial component with a timeline at the bottom, which has a two-fold purpose that we will describe in more detail.

Timeline

On the bottom, just below the radial component, we have included a dot plot (timeline) of the events that occur in the network (see Fig. 4-B). The timeline has a two-fold purpose: (i) it is used as an interaction tool that allows users to select custom ranges of time, individual events, or periods of time defined by historic events that occur. This means users can define their own temporal granularity to inspect the data; (ii) it is showing a different aspect of the data, aggregating the display of events for all related entities that persist in the network (people, locations, etc.), and showing patterns and trends that are associated with the historical development of the City of Vienna. It provides context to the data by embedding it in different historical frames of reference. Historic events are important for our application domain and expert users as they show a different perspective on the data. Historic events are represented as background areas along the timeline with a darker shading and have an icon that displays a tooltip with more detailed information about it. The intent of this component is also to guide users to more interesting periods of time and enable them to explore these in more detail (related to *T1: Present, T3: Explore, and T4: Identify*).

Temporal zoom

The temporal zoom interaction was conceived as a way to mitigate overlaps and very dense areas of the visualization by being able to select a span of time and have the granularity of that time span extended to the full size of the circle (see Figs. 4, 5, 6, and 7). It is also a way to support in-depth exploration and can provide different overviews depending on the selection of the user (*T1: Present* and *T3: Explore*). This interaction technique allows us to give more attention and drawing area to the selected period of time, making the data points falling into that time span spread out and declutter (related to *T3: Explore*).

Temporal inversion

Similarly, we have also explored how different approaches could be combined together to mitigate very dense areas, and overlaps created by busy periods of time. This is an issue with radial visualizations as points closer to the center of the circle receive less drawing area compared to those that are placed at the boundary of the circle. This could potentially affect the perceived importance of objects and lead the users to wrong insights and hypotheses or skew the importance of entities in an unintended way. We experimented with inverting the direction of time and allowing the user to interactively change this setting depending on the period of time that is of interest (see Fig. 8) (related to *T1: Present and T3: Explore*).

Ordering & grouping

Combinations of ordering and grouping can be used to detect similarities, differences, or patterns in event-based networks. Grouping can be used to position persistent entities, such as people, organizations, or locations, with similar properties in the same sector and can be further combined with ordering criteria that would reorder the rays in each sector according to some metric or criteria. Ordering and grouping can be used separately or combined to explore and investigate patterns that emerge under certain configurations or to explore if different groupings exhibit more similarities or more differences (related to *T2: Compare and T4: Identify*).

Fig. 9-A, B, and C show three different ways to order and group people, and how they affect the visual subject. In each case, different trends stand out. In Fig. 9-A people are ordered by their birth date and grouped by their role (composer, conductor, musician). A spiral forms growing outward, which is complementary to that of Fig. 9-B, where they are ordered by their death date and grouped by role. In the latter, the spiral signifies a ceiling, the limit of one's life. It is much easier to identify people that are alive in the second case, as all people that are still alive will form a contiguous arc. From these groupings and orderings we can see the average period of activity for the different roles and when this is put into historical context it can show shifts in musical motifs and trends. Additionally, the year 1945 (denoted by the dark gray circle) that is used by our domain experts as a baseline for comparison separates the data and we can see the difference between the types of events and their frequency before and after the end of World War II. Another insight that can be gained from this kind of grouping and ordering is that earlier on the bulk of the events related to people happened after their death as opposed to later periods of time. This implies that after the end of World War II the City of Vienna started honoring musicians during their lifetime as opposed to after their death in a pursuit of developing its public image as a "Music City". It is important to note that the data remains the same in all cases. By re-ordering and grouping them according to different criteria it is possible to obtain unique patterns that can tell us something about the underlying behavior of the data. This can also be used for other entities, for example users can also explore the subject of locations, their occupancy, and the change of their thematic nature over time by using different grouping and ordering options. Another use case could be exploring organizations and what different events they have contributed to over the course of time relative to their political affiliation.

Details

As mentioned in Section 2.3, we also would like to facilitate in-depth exploration and support verification of data points (people, locations, events, etc.). In our approach we achieve that by providing detail panels, where the users can explore the particularities of different entities, lookup their characteristics, and their relationships to other objects in the data. Our main objective with the detail panels was to engage users in learning new things and was meant as a way of answering questions that arouse during exploration and validation of the data. At the same time we wanted the detail panel to provide an informative and descriptive interface for users to access the details of objects and explore related entities, including aspects of the data that are not visualized, such as descriptive texts and material sources (related to *T3: Explore and T5: Verify*). When an object is selected the detail panel opens up and the radial visualization is reconfigured to focus and highlight the current selection, which enables the exploration of the temporal context and provides details.

5. Evaluation

In this section we describe the methodology employed in the evaluation of our approach, we present the procedure we followed, discuss the results of our interviews, and summarize our findings highlighting the lessons learned along the way.

5.1. Procedure: Evaluation setting

Our evaluation was designed mainly by domain experts in the spirit of a history exam, to be performed with the aid of Circular. In this regard, we would like to remark that the majority of tasks are phrased as questions rather than tasks in the traditional

Fig. 10. The results of the user study. The bars represent the level of difficulty the participants experienced for each task category as a percentage (0%–100%). The categories E (easy); M (medium); H (hard) represent a qualitative ranking of the performance of the participants for the questions from our evaluation. For the different tasks in the evaluation we can observe that our approach supports most tasks well (E), with some minor issues (M), and severe difficulties (H) were very rarely encountered.

sense of VA research. Essentially, the participants were asked to present their answers for different questions by using our approach and to compare, explore, and identify particular elements of the visualization that were the focal point of the user study.

We conducted a small-scale user study with both domain experts and VA experts, that was structured as a task-based evaluation. Throughout the study we encouraged the participants to think aloud (Lewis, 1982) as they performed tasks and to voice their thoughts and reasoning process along with feedback about interacting with our prototype.

We conducted the user study with six participants. Three of the participants were domain experts with little experience using visualizations to solve tasks, but possessing domain specific knowledge about the data. The other half of the participants were visualization experts with a more technical background but lacking knowledge in this specific application domain and the data itself. The sessions lasted around 60 min and were structured as follows: (i) onboarding and introduction to our approach; (ii) letting the users interact with and get used to our approach by performing some exploratory tasks; (iii) task-based evaluation; (iv) general feedback. There were a total of 18 different tasks belonging to five task categories (see Section 2.3 for more information and examples on the tasks).

As our approach is situated at the intersection of multiple disciplines, it is not trivial to formally evaluate it. Due to the fact that most of the tasks are not analytical, it becomes even more difficult to decide on an evaluation methodology as tasks are not a well-defined subject (Schetinger et al., 2019). Common metrics that are used to formally evaluate and prove the efficacy of approaches in well-established and defined disciplines are difficult to apply to domains such as digital humanities. We believe that metrics, such as, measuring error rates or time per task are not relevant for determining the usefulness of such approaches. Furthermore, our intent was not to conduct a comparative evaluation with other representations of the same data (i.e., rectilinear timeline visualizations or other traditional network layout representations) as these were not in line with our evaluation interests. Our main aim was to investigate the capability of radial representations for representing event-based networks and to assess how domain and visualization experts perceive it in order to identify the benefits and drawbacks of our approach. For this purpose we investigated alternative evaluation methodologies, including the work of Klein et al. on sense-making models (Klein et al., 2006a,b) and arrived to the conclusion that we would like to evaluate how well our approach supports the tasks outlined in Section 2.3.

In our evaluation we used a simple qualitative ranking to identify issues that made it difficult for users to perform their tasks. We encoded each of the participants answers as easy (E), medium (M), or hard (H). If the answering process was straightforward and the user could translate the question into a set of interactions, perform them, and identify the answer, this counts as an easy (E) answer. If the user had to try different approaches to reach the goal or there was something that caused confusion or was unclear, it counts as medium (M). If the user was stuck and the evaluators had to intervene and provide guidance, it is counted as hard (H). The encoding of the answers was done during each session by two visualization experts. After each session there was a short discussion about the results of the participant to reach a final decision on the encoding. Additionally, we had a domain expert verifying the correctness of the participants' answers and clarifying any misconceptions about the data. There was little ambiguity observed when deciding between which of the categories to register, and the gathered results showed consistency among the different questions.

5.2. Discussion

In this section we will discuss how each of the tasks was solved in the user-study along with the participants feedback and personal opinions. We consider both the results of the domain experts and visualization experts. A general overview of how our approach performed for the different task categories can be seen in Fig. 10.

T1: Present: For the most part our approach provides a good overview of the data and allows for different ways to present the historical narratives, data, and specific details about different entities. In this task category we noticed that most participants managed to answer the questions with little to no difficulties, apart from two questions that were outliers, so we inspected these in more detail (*Q3 - Which themes are related to the events taking place at Stadthalle?* and *Q12 - In which district(s) did fewest Events take place?*) to determine the cause of the issues (see Fig. 10-T1). We expected the participants to present their answers to the questions in this task category by interacting with different components of our approach, specifically, the operations to filter, order and group data points. Through navigation methods, such as zooming, panning, and triggering on-demand detail views, we expected the participants to present the elements of interest in the visualization. We anticipated that the answers to the questions would become evident in the radial and timeline components of our approach. For the first question (Q3) the participants needed to identify the themes related to events happening at a specific location and then enumerate these. We observed that this turned out to be difficult for most users due to a usability issue. When selecting a location, the visualization is reconfigured to highlight the events taking place at that location and hides interface elements, this in turn made it difficult for our users to identify exactly which themes were related to the events, some users attempted to access this information from the detail panel of each event and keep track of this information whereas some answered the question by listing the colors they observed along the ray of this location. The second question (Q12) involved altering the state of the visualization to display a different aspect of the data, specifically, grouping locations by district and then quantifying the amount of events taking place in each grouping in order to gain an overview of which district had the most events organized.

Our findings are that these two questions were fairly complex as they involved keeping track of multiple details and associations throughout different views and states, which lead some participants to confusion and affected their performance at this task. Overall, the evaluation results lead us to believe that our approach supports *T1: Present* because the majority of the questions were answered with ease. An additional observation we made is that some tasks were easier for our domain experts to accomplish because they had the advantage of possessing more in-depth domain knowledge so they could easily identify this information in the visualization and confirm what they already knew. Furthermore, we noticed that the domain experts spent more time exploring the different options and settings that Circular offers to reconfigure the visualization and got acquainted to it quite fast.

T2: Compare: For the second task category the results of the evaluation lead us to believe that our approach performs better at compare tasks than present tasks (see Fig. 10-T2). Overall, the participants answered most questions with ease apart from some difficulties that we investigated more carefully. Our findings are that the main problem stems from the same two questions as the previous task (Q3 and Q12), where the goal is to compare quantitative information of groupings and identify the details of thematic changes over time and list these. The issues identified here were related to elements of the user interface being hidden during selections of data points and the unfamiliarity with the set of options for grouping and ordering criteria we offer. This caused the participants to have to keep track of many details throughout the tasks in this category and increased their cognitive load. The domain experts did perform overall better and managed to answer the questions with no difficulty due to the fact that they invested more time, initially, to explore the different options and settings that are used to reconfigure the state of the visualization. We expected the participants to answer the questions in this task category by interacting with the different options in our approach for grouping and ordering data and also combining these to modify the state of the visualization. The use of different arrangement methods to reorganize elements of the visualization spatially should support the users in the comparison task. The aim of the participants is to configure the radial component in such a way as to better outline differences between categories of the data in order to detect changes that occur over time and compare these. In the general case, we discovered some limitations that impacted how users performed at questions including counting and grouping tasks. Overall, the results of our evaluation for task *T2: Compare* show that identifying objects of similar characteristics and comparing them is supported in Circular and users are able to successfully complete these tasks. The difficulties we discovered in the evaluation were related to counting and comparing the quantitative information of different groups of entities as this information had to be manually acquired by the participants, but the relative sizes of the sectors managed to convey a sense of this to the participants. As an improvement we believe it would be best to either encode the result of the groupings and count of specific properties straight into the visualization itself (e.g., by having the numbers and percentages displayed in each grouping) or offer an additional view dedicated to the display of such information.

T3: Explore: The intent of this task was to provide support for accessing the details of each object, including any material sources that are related to it. These sources were gathered by our domain experts and were used to shape and structure the data. For our domain experts, it was very important to be able to see and explore the sources and various details that were associated with each object, including related entities. In our approach we provide a few different ways to explore the details of the data: (i) a detail panel that opens up when an object is selected and is configurable for multiple object types as they have different properties associated with different importance; (ii) a list view that users can use to explore a different perspective that depicts a more high-level representation of the data. For this task category we expected the participants of the study to find the answers of the respective questions by filtering according to categorical attributes and navigating to the detail view in order to access and explore further information, including material sources (video and photographic material, audio recordings, and textual resources). By selecting elements of the visualization or by brushing specific periods of time to explore, the users should be able to explore the properties and relationships of elements and discover anomalous behavior (e.g., discontinuity of periodic events). We also expected more participants would make use of the list view that provides a high-level overview of the data that is currently selected, but this was not the case as most participants preferred using the radial component and detail panel to answer the questions.

For this category the results of our evaluation and user feedback show that our approach supports the *T3: Explore* task very well and that the participants felt that the solution provides suitable interactions for examining the details and intricacies of the complex historical data we are depicting. The difficulties (M) that can be seen in our findings (see Fig. 10-T3) stem from the fact that in our implementation, when selecting individual data points, the visualization is reconfigured to highlight the selected data and the rest of the information fades away, which was a bit confusing for most participants. While, initially, we felt this was a playful interaction made to capture the participants' attention and interest, it seemed to be counter-intuitive for most. Overall, we observed that there were very few difficulties encountered in the questions belonging to this task category and most of them could be performed in a straightforward fashion.

T4: Identify: One of the major requirements (see Section 2.3) when designing and developing our approach was to enable users to see patterns and trends throughout the City of Vienna's history, including continuities and discontinuities of certain themes over time and how they are related to different people and/or locations. For the majority of questions in this task category the participants managed to provide answers with little to no difficulties and this is also reflected in the results of our study (see Fig. 10-T4). We expected the participants to identify the persistent entities of interest (people or locations) by using different combinations of grouping and ordering options and observing their temporal developments (related events and themes over time). The goal for this task category is for the participants to identify targets or sets of targets that are to be further investigated to drill-down to their properties and relationships. Identifying outliers or data points that exhibited interesting temporal developments are also important in this context and we anticipated this to be visible under certain configurations of the visualization. We expected the answer to each question in this task category to be evident in the radial component of our approach. The issues that the participants had with the tasks in this category were related to the questions, that were also associated with difficulties in the other task categories. The reason for these difficulties arose from the interactions to access the details of specific data points. These interactions enable the detail panel and hide legends and other parts of the visualization that are not part of the current selection, but provide additional information that is necessary to answer the questions. The participants expressed some frustration about the disappearing elements and this can be also observed with the increase in difficulty of the questions in this category (increase in the M and H categories relative to other tasks, see Fig. 10-T4).

Overall, the results from the evaluation indicate that our approach supports this task category fairly well despite some usability issues. Most questions were easily answered, but there were some difficulties with the more complex questions, requiring a longer path of interactions and keeping track of interim details to provide an answer. We believe that by improving some of the usability issues and making certain details more explicit we could achieve better support for the tasks in this category.

T5: Verify: The data that we used in our approach is modeled as an event-based network, shaped and scoped by our domain experts. The domain experts continuously modify data points and entities to better frame the data and their perspective. Our dataset is the result of historical research conducted by the experts, sampling key points and important figures, establishing narratives, and identifying the most prominent themes in the City of Vienna's history. For this purpose, it was very important to also provide support for validating and verifying the data that was entered and visualized. For our domain experts this can be a cumbersome process as the data is manually entered and curated. Since most database solutions usually require a level of technical expertise and knowledge, validating the data and finding errors becomes a bit of a challenge. Furthermore, most database solutions generally do not provide any visual means to gain an overview and contextualize the data. Whereas when presenting the data in a visual manner to our domain experts the erroneous data became immediately obvious to them and supported them resolving the inaccuracies and inconsistencies of the data. In this task category we expected the participants to query the visualization for particular persistent entities based on their characteristics and utilize the radial and detail components of our approach to provide answers related to specific properties of those entities (i.e., dates, names, and relationships). We constructed the questions in this task category as a way to validate and verify the correctness of the data that is being depicted in the visualization. From the results of the evaluation and feedback from the domain experts we can also conclude that Circular provides support for the tasks associated with *T5: Verify*. Any issues with the data immediately become obvious in the visualization and through the use of detail panels and interactions we offer users multiple views to identify and correct this information. In Fig. 10-T5 we notice that there were no difficulties encountered throughout the tasks in this category by the participants, but during the sessions we observed that there were some attempts at trying different approaches and paths to reach the answer for the same question and confirm the validity of the data being inspected, which is expected.

The tasks presented in this evaluation are low-level tasks, but we also consider the different ways that they can be combined together to form more complex paths of interaction and to support higher level tasks as well. In this regard we would like to draw parallels to the synoptic vs. elementary tasks introduced by Andrienko and Andrienko (2006). This notion is further supported by the multi-level task typology we employ in our research, in that sensemaking is built upon combinations of such low-level tasks (Brehmer and Munzner, 2013).

5.3. Feedback

All participants agreed that this was an interesting approach to explore linked historical data and offered an engaging interface to delve into the details of events, people, locations, and themes seeking changes over time. Specific attention was brought to the radial component as the participants stated that this was an aesthetically pleasing, compact, and interesting way of encoding this type of information. These statements also align with our initial assumptions that a radial visualization would be an engaging interface, which is also backed by existing literature (Borkin et al., 2013; Hohman et al., 2020) (see Section 4.1). One of the domain experts mentioned that radial visualizations were something new to them and they were an interesting way of representing data. The interaction techniques that we designed and implemented were appreciated and fostered user engagement and drove the participants' curiosity to explore the data in more detail and discover new information. We observed that even after the evaluation session had concluded some participants spent additional time exploring the data and discussing topics that were of interest to them. Overall, the results of our evaluation lead us to believe that our approach achieves an interesting way of engaging users that would like to explore linked historical data and discover new insights.

According to the participants our solution provides a nice and compact overview of historical data. The general consensus is that it succeeds at representing both the relational and temporal information between disparate entities using appropriate metaphors and it enabled for exploration in multiple dimensions.

The participants also commented on the ability of our approach to show the progression and evolution of themes and patterns over space and time. One of the visualization experts that participated stated that our node-sizing could be considered a double-edged blade as increasing a node's size that is positioned closer to the circle's center would presumably cause even more overlaps, but it was also appreciated and considered interesting as the user could interactively change it based on their interests and explore events that are more connected to themes, material sources, people, or other types of entities.

There were also some drawbacks to our approach, including that it has a bit of a learning curve and some onboarding (Stoiber et al., 2019) or guidance (Ceneda et al., 2017) would be useful to get the participants better acquainted with the visualization and guide them towards interesting points in history. Furthermore, we identified minor usability issues throughout the evaluation, which made some of the questions in the study difficult to answer. In the current state of our approach interface elements are hidden when selecting data points to explore their details, which in turn caused difficulties for the participants of our study. Initially, this was intended as an engaging way to keep the focus of the users on their selection, but according to the results and feedback of the evaluation, this was not considered helpful or useful. The most notable functionality that was missing was the ability to navigate back and forth between different selections in the detail view alongside a history of interactions so that users could easily go to previous states of the visualization. Additionally, the participants suggested that we provide the ability to lock the temporal selection of the user when interacting with the different components and displaying a list of the events that are happening within that selection along with the people, locations, and themes they are related to. These features would be relevant for future work as quite some participants expressed that it would increase the usability of our approach and improve the ability to navigate the data. There were also a few comments about the data quality, completeness, and uncertainty, but these are out of the scope of our evaluation.

6. Conclusion and future work

In this paper we presented our approach, Circular, for visualizing event-based networks in the domain of digital humanities. It is an interactive exploration environment that utilizes a radial layout to represent a large amount of linked historical research modeled as an event-based network that is embedded in spatial and temporal frames of reference. In Circular we put the spotlight on the temporal and relational aspects of the data and visualize event-based networks in a compact, aesthetic, and engaging manner. We provide users with multiple ways to slice the network based on their tasks and interests and also explore how grouping and ordering combinations change the network's topology along with what insights can be extracted from the visualization. We have evaluated our approach by conducting a small-scale user study with both domain experts and visualization experts that highlights the usefulness and uniqueness of our approach but also its limitations and drawbacks. For future work, other than improving upon the current limitations and usability issues that were identified, we would like to augment the topology of the network with more information, such as historic events and sources that add an additional layer of complexity to the network but also provide more context and narrative to the data. Additionally, we believe implementing semantic zooming would be an interesting approach to control the information density based on the granularity specified by the user and offer a solution to the clutter and overlaps close to the circles center.

CRediT authorship contribution statement

Velitchko Filipov: Conceptualization, Methodology, Visualization, Software, Writing – Original draft preparation, Writing - review & editing, Investigation, Validation. **Victor Schetinger:** Methodology, Visualization, Supervision, Writing - Original draft preparation, Writing - review & editing, Investigation, Validation. **Kathrin Raminger:** Data curation, Resources, Investigation, Validation. **Nathalie Soursos:** Data curation, Resources, Writing – Original draft preparation. **Susana Zapke:** Data curation, Resources, Project Administration, Funding acquisition. **Silvia Miksch:** Supervision, Writing - Reviewing & Editing, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was conducted within the framework of the project "Interactive Music Mapping Vienna" (AR384-G24) and "Knowledge–Assisted Visual Analytics" (P31419-N31) funded by the Austrian Science Fund (FWF).

References

Ahn, J.-w., Plaisant, C., Shneiderman, B., 2014. A task taxonomy for network evolution analysis. IEEE TVCG 20 (3), 365–376. http://dx.doi.org/10.1109/TVCG.2013.238.

Alsallakh, B., Aigner, W., Miksch, S., Hauser, H., 2013. Radial sets: Interactive visual analysis of large overlapping sets. IEEE TVCG 19 (12), 2496–2505. http://dx.doi.org/10.1109/TVCG.2013.184.

Andrienko, N., Andrienko, G., 2006. Exploratory Analysis of Spatial and Temporal Data: A Systematic Approach. Springer, Berlin, http://dx.doi.org/10.1007/3-540-31190-4.

Bale, K., Chapman, P., Barraclough, N., Purdy, J., Aydin, N., Dark, P., 2007. Kaleidomaps: A new technique for the visualization of multivariate time-series data. Inf. Vis. 6 (2), 155–167. http://dx.doi.org/10.1057/palgrave.ivs.9500154.

Beck, F., Burch, M., Diehl, S., Weiskopf, D., 2014. The state of the art in visualizing dynamic graphs. In: Borgo, R., Maciejewski, R., Viola, I. (Eds.), EuroVis - STARs. The Eurographics Association, http://dx.doi.org/10.2312/eurovisstar.20141174.

Borkin, M., Vo, A.A., Bylinskii, Z., Isola, P., Sunkavalli, S., Oliva, A., Pfister, H., 2013. What makes a visualization memorable? IEEE TVCG 19 (12), 2306–2315. http://dx.doi.org/10.1109/TVCG.2013.234.

Börner, K., Eide, O., Mchedlidze, T., Rehbein, M., Scheuermann, G., 2019. Network visualization in the humanities (Dagstuhl seminar 18482). In: Börner, K., Eide, O., Mchedlidze, T., Rehbein, M., Scheuermann, G. (Eds.), Dagstuhl Rep. 8 (11), 139–153. http://dx.doi.org/10.4230/DagRep.8.11.139.

Brehmer, M., Lee, B., Bach, B., Riche, N.H., Munzner, T., 2017. Timelines revisited: A design space and considerations for expressive storytelling. IEEE TVCG 23 (9), 2151–2164. http://dx.doi.org/10.1109/TVCG.2016.2614803.

Brehmer, M., Munzner, T., 2013. A multi-level typology of abstract visualization tasks. IEEE TVCG 19 (12), 2376–2385. http://dx.doi.org/10.1109/TVCG.2013.124.

Burch, M., Diehl, S., 2008. TimeRadarTrees: VIsualizing dynamic compound digraphs. Comput. Graph. Forum 27, 823–830. http://dx.doi.org/10.1111/j.1467-8659.2008.01213.x.

Burch, M., Weiskopf, D., 2014. On the benefits and drawbacks of radial diagrams. In: Handbook of Human Centric Visualization. Springer New York, pp. 429–451. http://dx.doi.org/10.1007/978-1-4614-7485-2_17.

Cappers, B.C.M., van Wijk, J.J., 2017. Exploring multivariate event sequences using rules, aggregations, and selections. IEEE TVCG 24, 532–541. http://dx.doi.org/10.1109/TVCG.2017.2745278.

Castermans, T., Verbeek, K., Speckmann, B., Westenberg, M.A., Koopman, R., Wang, S., van den Berg, H., Betti, A., 2019. SolarView: Low distortion radial embedding with a focus. IEEE TVCG 25 (10), 2969–2982. http://dx.doi.org/10.1109/TVCG.2018.2865361.

Ceneda, D., Gschwandtner, T., May, T., Miksch, S., Schulz, H., Streit, M., Tominski, C., 2017. Characterizing guidance in visual analytics. IEEE TVCG 23 (1), 111–120. http://dx.doi.org/10.1109/TVCG.2016.2598468.

Cleveland, W.S., McGill, R., 1984. Graphical perception: Theory, experimentation, and application to the development of graphical methods. J. Amer. Statist. Assoc. 79 (387), 531–554. http://dx.doi.org/10.2307/2288400.

Dang, T.N., Pendar, N., Forbes, A.G., 2016. TimeArcs: Visualizing fluctuations in dynamic networks. Comput. Graph. Forum http://dx.doi.org/10.1111/cgf.12882.

Di Bartolomeo, S., Pandey, A., Leventidis, A., Saffo, D., Syeda, U.H., Carstensdottir, E., Seif El-Nasr, M., Borkin, M.A., Dunne, C., 2020. Evaluating the effect of timeline shape on visualization task performance. In: Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems. CHI '20, Association for Computing Machinery, New York, NY, USA, pp. 1–12. http://dx.doi.org/10.1145/3313831.3376237.

Diehl, S., Beck, F., Burch, M., 2010. Uncovering strengths and weaknesses of radial visualizations—an empirical approach. IEEE TVCG 16 (6), 935–942. http://dx.doi.org/10.1109/TVCG.2010.209.

Draper, G., Livnat, Y., Riesenfeld, R.F., 2009. A survey of radial methods for information visualization. IEEE TVCG 15 (5), 759–776.

Du, F., Shneiderman, B., Plaisant, C., Malik, S., Perer, A., 2017. Coping with volume and variety in temporal event sequences: Strategies for sharpening analytic focus. IEEE TVCG 23 (6), 1636–1649. http://dx.doi.org/10.1109/TVCG.2016.2539960.

Elzen, S., Holten, D., Blaas, J., van Wijk, J.J., 2014. Dynamic network visualization withextended massive sequence views. IEEE TVCG 20 (8), 1087–1099. http://dx.doi.org/10.1109/TVCG.2013.263.

Filipov, V., Soursos, N., Schetinger, V., Zapke, S., Miksch, S., 2019. Exiled but not forgotten: Investigating commemoration of musicians in vienna after 1945 through visual analytics. In: 3rd Conference on Biographical Data in a Digital World. http://dx.doi.org/10.17605/OSF.IO/RVHDS.

Hall, K.W., Bradley, A.J., Hinrichs, U., Huron, S., Wood, J., Collins, C., Carpendale, S., 2020. Design by immersion: A transdisciplinary approach to problem-driven visualizations. IEEE TVCG 26 (1), 109–118. http://dx.doi.org/10.1109/TVCG.2019.2934790.

Hinrichs, U., El-Assady, M., Bradley, A.J., Forlini, S., Collins, C., Risk the drift! Stretching disciplinary boundaries through critical collaborations between the humanities and visualization. In: 2nd Workshop on Visualization for the Digital Humanities.

Hohman, F., Conlen, M., Heer, J., Chau, D.H.P., 2020. Communicating with interactive articles. Distill http://dx.doi.org/10.23915/distill.00028, URL https://distill.pub/2020/communicating-with-interactive-articles.

Keim, D.A., Schneidewind, J., Sips, M., 2004. Circleview: A new approach for visualizing time-related multidimensional data sets. In: Proceedings of the Working Conference on Advanced Visual Interfaces. AVI '04, Association for Computing Machinery, New York, NY, USA, pp. 179–182. http://dx.doi.org/10.1145/989863.989891.

Klein, G., Moon, B., Hoffman, R.R., 2006a. Making sense of sensemaking 1: Alternative perspectives. IEEE Intell. Syst. 21 (4), 70–73. http://dx.doi.org/10.1109/MIS.2006.75.

Klein, G., Moon, B., Hoffman, R.R., 2006b. Making sense of sensemaking 2: A macrocognitive model. IEEE Intell. Syst. 21 (5), 88–92. http://dx.doi.org/10.1109/MIS.2006.100.

Krzywinski, M., Schein, J., Birol, I., Connors, J., Gascoyne, R., Horsman, D., Jones, S.J., Marra, M.A., 2009. Circos: An information aesthetic for comparative genomics. Genome Res. 19 (9), 1639–1645. http://dx.doi.org/10.1101/gr.092759.109.

Lamqaddam, H., Brosens, K., Truyen, F., Beerens, R.J., De Prekel, I., Verbert, K., 2018. When the tech kids are running too fast: Data visualisation through the lens of art history research. IEEE TVCG.

Lee, B., Plaisant, C., Parr, C.S., Fekete, J.-D., Henry, N., 2006. Task taxonomy for graph visualization. In: Proceedings of the 2006 AVI Workshop on beyond Time and Errors: Novel Evaluation Methods for Information Visualization. BELIV '06, Association for Computing Machinery, New York, NY, USA, pp. 1–5. http://dx.doi.org/10.1145/1168149.1168168.

Lewis, C., 1982. Using the Thinking-Aloud Method in Cognitive Interface Design. IBM TJ Watson Research Center Yorktown Heights, NY, http://dx.doi.org/10.1080/00140139008927157.

Lincoln, M., 2016. Social network centralization dynamics in print production in the low countries, 1550-1750. Int. J. Digit. Art Hist. (2), http://dx.doi.org/10.11588/dah.2016.2.25337.

Liu, Q., Hu, Y., Shi, L., Mu, X., Zhang, Y., Tang, J., 2015. EgoNetCloud: EVent-based egocentric dynamic network visualization. In: 2015 IEEE Conference on Visual Analytics Science and Technology (VAST). pp. 65–72. http://dx.doi.org/10.1109/VAST.2015.7347632.

Maçãs, C., Machado, P., 2018. Radial calendar of consumption. In: 2018 22nd International Conference Information Visualisation (IV), pp. 96–102.

Mashima, D., Kobourov, S., Hu, Y., 2012. Visualizing dynamic data with maps. IEEE TVCG 18 (9), 1424–1437. http://dx.doi.org/10.1109/TVCG.2011.288.

Miksch, S., Aigner, W., 2014. A matter of time: Applying a data–users–tasks design triangle to visual analytics of time-oriented data. Comput. Graph. 38, 286–290.

Monroe, M., Lan, R., Lee, H., Plaisant, C., Shneiderman, B., 2013. Temporal event sequence simplification. IEEE TVCG 19 (12), 2227–2236. http://dx.doi.org/10.1109/TVCG.2013.200.

Monroe, M., Lan, R., Morales del Olmo, J., Shneiderman, B., Plaisant, C., Millstein, J., 2013. The challenges of specifying intervals and absences in temporal queries: A graphical language approach. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. CHI '13, Association for Computing Machinery, New York, NY, USA, pp. 2349–2358. http://dx.doi.org/10.1145/2470654.2481325.

Munzner, T., 2009. A nested model for visualization design and validation. IEEE TVCG 15 (6), 921–928. http://dx.doi.org/10.1109/TVCG.2009.111.

Munzner, T., 2014. Visualization Analysis and Design. A K Peters Visualization Series, CRC Press.

Nightingale, F., 1987. Notes on Matters Affecting the Health, Efficiency, and Hospital Administration of the British Army: Founded Chiefly on the Experience of the Late War. Harrison and Sons, St. Martin's Lane, WC.

O'Madadhain, J., Hutchins, J., Smyth, P., 2005. Prediction and ranking algorithms for event-based network data. ACM SIGKDD Explor. Newslett. 7 (2), 23–30. http://dx.doi.org/10.1145/1117454.1117458.

Plaisant, C., Milash, B., Rose, A., Widoff, S., Shneiderman, B., 1996. Lifelines: Visualizing personal histories. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. CHI '96, Association for Computing Machinery, New York, NY, USA, pp. 221–227. http://dx.doi.org/10.1145/238386.238493.

Playfair, W., 1801. The Statistical Breviary: Shewing, on a Principle Entirely New, the Resources of Every State and Kingdom in Europe; Illustrated with Stained Copper-plate Charts the Physical Powers of Each Distinct Nation with Ease and Perspicuity: to which is Added, a Similar Exhibition of the Ruling Powers of Hindoostan. T. Bensley, Bolt Court, Fleet Street.

Schetinger, V., Raminger, K., Filipov, V., Soursos, N., Zapke, S., Miksch, S., 2019. Bridging the gap between visual analytics and digital humanities: Beyond the data-users-tasks design triangle. In: Workshop on Visualization for the Digital Humanities, IEEE VIS 2019. http://dx.doi.org/10.6084/m9.figshare.12174483.

Schich, M., Song, C., Ahn, Y.-Y., Mirsky, A., Martino, M., Barabási, A.-L., Helbing, D., 2014. A network framework of cultural history. Science 345 (6196), 558–562. http://dx.doi.org/10.1126/science.1240064.

Schulz, H., Nocke, T., Heitzler, M., Schumann, H., 2013. A design space of visualization tasks. IEEE TVCG 19 (12), 2366–2375. http://dx.doi.org/10.1109/TVCG.2013.120.

Shi, Y., Zhao, Y., Zhou, F., Shi, R., Zhang, Y., 2018. A novel radial visualization of intrusion detection alerts. IEEE Comput. Graph. Appl. 38 (6), 83–95. http://dx.doi.org/10.1109/MCG.2018.2879067.

Simonetto, P., Archambault, D., Kobourov, S.G., 2020. Event-based dynamic graph visualisation. In: IEEE TVCG, pp. 2373–2386.

Stoiber, C., Grassinger, F., Pohl, M., Stitz, H., Streit, M., Aigner, W., 2019. Visualization onboarding: Learning how to read and use visualizations. In: IEEE Workshop on Visualization for Communication. IEEE, Vancouver, BC, Canada, http://dx.doi.org/10.31219/osf.io/c38ab.

Tory, M., Moller, T., 2004. Rethinking visualization: A high-level taxonomy. In: IEEE Symposium on Information Visualization. pp. 151–158. http://dx.doi.org/10.1109/INFVIS.2004.59.

Vehlow, C., Beck, F., Weiskopf, D., 2015. The state of the art in visualizing group structures in graphs. In: Eurographics Conference on Visualization (EuroVis) - STARs. The Eurographics Association, pp. 21–40. http://dx.doi.org/10.1111/cgf.12872.

Waldner, M., Diehl, A., Gracanin, D., Splechtna, R., Delrieux, C., Matkovic, K., 2019. A comparison of radial and linear charts for visualizing daily patterns. IEEE TVCG 26 (1), 1033–1042. http://dx.doi.org/10.1109/TVCG.2019.2934784.

Wertheimer, M., 1923. Untersuchungen zur Lehre von der Gestalt. II. Psychol. Res. 4 (1), 301–350.