EuroVis Workshop on Visual Analytics (2021)
J. Bernard and K. Vrotsou (Editors)

LFPeers: Temporal Similarity Search in Covid-19 Data

J. Burmeister[1] and J. Bernard[2] and J. Kohlhammer[1,3]

[1]Fraunhofer IGD, Germany, [2]University of Zurich, Switzerland [3]TU Darmstadt, Germany

Figure 1: *The search interface of LFPeers, displaying an ongoing search for countries similar to the USA using three similarity criteria. Orange indicates the set of currently included similar countries, grey the global data as context. For the active criteria the user has defined the temporal attribute "Stringency Index" and a time range (marked in grey), causing the table of countries to be sorted by similarity to the reference. A t-SNE projection and world map provide further insights into the geospatial and structural nature of the selection.*

Abstract

While there is a wide variety of visualizations and dashboards to help understand the data of the Covid-19 pandemic, hardly any of these support important analytical tasks, especially of temporal attributes. In this paper, we introduce a general concept for the analysis of temporal and multimodal data and the system LFPeers that applies this concept to the analysis of countries in a Covid-19 dataset. Our concept divides the analysis in two phases: a search phase to find the most similar objects to a target object before a time point t_0, and an exploration phase to analyze this subset of objects after t_0. LFPeers targets epidemiologists and the public who want to learn from the Covid-19 pandemic and distinguish successful and ineffective measures.

1. Introduction

Since the first broad wave of infections in early 2020, the Covid-19 pandemic is having a profound effect on almost all aspects of personal life, business, and government around the world. While all countries had to contain the same virus (and its mutations), there was a wide variety of policy measures, approaches, and outcomes in terms of the number of infections, deaths, and reproduction rates. This variety is not solely rooted in the political system or the economical situation of a country, as even quite similar countries in the EU had to endure more or less severe situations. Many researchers around the world have sought to understand Covid-19 indicators with various analyses, infographics, and information visualization techniques. Many interactive dashboards have emerged in short time [Com20], including prominent Covid-19 trackers by the Johns Hopkins University [DDG20] or the Financial Times [Tim20] to inform the general public and to consult governments all over the world. With an emphasis on temporal analysis, many approaches show *how* situations are like and *how* measures have developed over time. However, less approaches also help explaining *why* situ-

DOI: 10.2312/eurova.20211098

ations have emerged (cause-and-effect), and even fewer approaches allow to infer about what we can learn from past developments of individual countries (learning from peers). The complexity of data types and analysis tasks for these envisioned goals calls for visual analytics solutions.

This paper focuses on the analysis of cause-and-effect within a combined data set of temporal [AMST11] and multimodal [KH13] attributes, collected by Google [W*20] during the Covid-19 pandemic until today. To learn from *peers*, our approach includes user-definable nearest neighbor search operations in early temporal phases, in combination with the exploration of the development of peers in later phases. Related to that, Afzal et al. provide a solution based on geographical visualization to assist in building public response plans [AGJ*20]. There are also several epidemiological approaches that address emergent pandemia [PBK*19] or visualizations related to epidemic management [YDH*17, AKMR16]. The work by Leite et al. [LSC*20] is related, with their goal to support the understanding of the spread of the disease in different scenarios also in relation to different government response events, while differing in their focus on predictive modeling. Looking beyond epidemiology, there have been several approaches targeting temporal data and similarity [WGGP*11, BSH*16, ZJGK10]. In a sense, we are extending earlier work by Cibulski et al. who established counterparts for the similarity of temporal observations and of corresponding output progressions [CMP*19].

Our principal idea is to analyze where and why countries sharing a similar situation started to deviate from each other. We build on the user-defined similarity of observations in a search phase (*before* t_0), followed by a detailed analysis of dissimilarities of temporal progressions in an exploration phase (*after* t_0). Users can also define the focus point t_0, as it is often exactly known when influencing factors (such as changes in policies) have been implemented, or interesting changes in observed measures may suggest focus points that are most useful. Pinpointing the policies and measures that were different from t_0 onward helps understand the causes and effects, and helps distinguish effective from irrelevant approaches or identifying irrelevant attributes for a perceived effect. None of the current Covid-19 analytical approaches known to us let users define a distance function and the similarity of countries and Covid-19 situations based on temporal and multimodal data. Our novel distinction between the time range before and after a focus point opens new avenues for time-oriented analysis of infection data on the per-country level. Further, while we were initially motivated by the Covid-19 analytics scenario, our concept takes a more general approach. Overall, we aim at the following contributions: first, we provide a novel concept for temporal and multimodal similarity search and exploration; second, we introduce and demonstrate a prototypical VA system, LFPeers (short for "learning from peers"), that implements this concept for Covid-19 data; third, we reflect on the current state of development and provide ideas for future work.

2. Concept

Our conceptual design approach is mainly based on the data-user-task triangle [MA14]. However, for the visualization and analysis design, we prioritize a clear and stringent task definition over an in-depth user characterization, as our primary aim is to open the analysis for broader yet undefined user groups.

2.1. Abstractions

We abstract from the Covid-19 domain and the dataset at hand, and define the analysis problem in a domain-agnostic way. We use the time axis to structure the presentation of these abstractions, as the evolution of the data over time nicely aligns with the order of steps in the workflow of analysts. Accordingly, we structure abstractions across the temporal domain of the dataset (see Figure 2).

Data abstraction Across the entire data analysis process, analysts focus on the very same type of data objects (here: countries) which are represented by multiple attributes. Mandatory for this type of approach is the existence of at least one time-dependent attribute, as the goal of our analyses is related to time-orientation. In Figure 2, time series data of all objects (for one focused attribute) is visualized with light gray color as a background layer. What makes the approach special is the existence of additional data attributes which may or may not develop over time. These additional data attributes will help to contextualize or even explain findings in the focused time-oriented attribute. It may be desirable to include different and multimodal sources of data into the approach [KH13], inspired by the idea to increase the variety of interesting connections [BSW*14] that can be drawn across attributes. In our case, we load data from various domains (epidemiology, government responses, weather, etc.) into LFPeers, all linked to the data object at hand (countries).

Problem statement Given a multimodal collection of data objects, the goal of analysts is to identify behavior of these objects over time and to explain variations of unexpected behavior to make informed decisions. Inspired by the idea to learn from peers and do better next time, most challenging aspects in the analysis workflow are to identify near-neighbor objects to compare these objects over time and to explain variation across these objects. The analysis problem is special, as analysts face the challenge of an interesting duality with respect to the temporal domain. In an early temporal phase of the time series, the challenge is to identify and select a subset of *most similar objects*. In contrast, in later phases of the time series the information need of analysts becomes exploratory [AA06]: it flips to the *identification and explanation of variation (dissimilarity)* within the previously selected subset. In Figure 2, a thin slope of orange time series indicates the identification and selection of data similar to the blue reference time series in the *search phase*. In contrast, the *exploration phase* after t_0 employs a fan metaphor to indicate that the subset of previously selected time series starts developing (micro) variations, which will be subject to in-depth analysis.

2.2. Tasks along the User Workflow

The workflow of analysts includes two phases. In an early *search phase* users retrieve nearest neighbors [BAP*05] prior to an end-point t_0 (see Figure 2). The *exploration phase* starts after the search phase, framed by a second endpoint t_1. In this phase, users identify and explain variances in the retrieved subset. As such, the approach adopts the exploratory search [WR09, Ber15] principle for time-oriented data [BWK*13] in a novel way. For the design of visual analytics support, we suggest the following set of abstracted tasks.

- **Search Phase**: retrieve nearest neighbors prior to t_0.
  - $\mathbf{T}_1$: select and analyze object of interest (here: country)
  - $\mathbf{T}_2$: define t_0 after the analysis of the object of interest

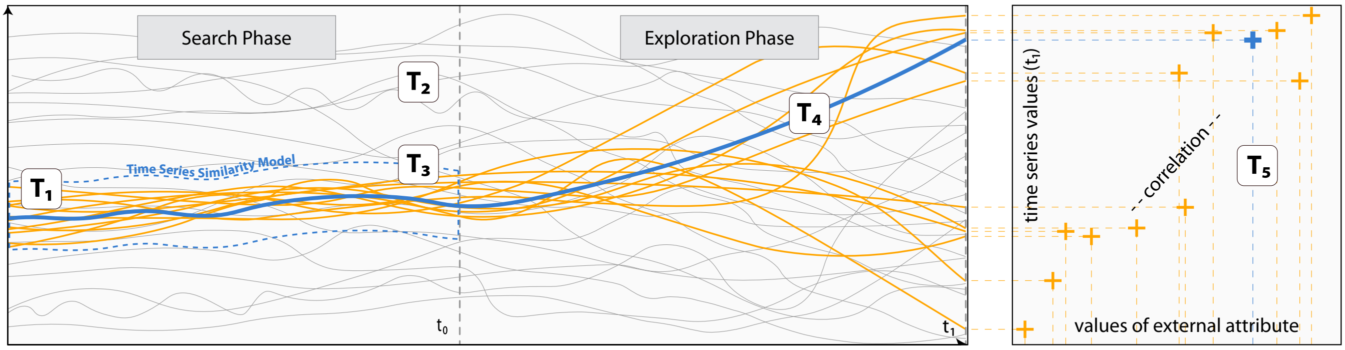

Figure 2: *Conceptual user workflow with five tasks that align with the temporal progression of the data. First, a user selects a focus object (**T**$_1$) represented by its time series (blue). Second, the user defines a temporal endpoint t_0 of the search phase (**T**$_2$). Third, based on the defined similarity, nearest neighbor items are automatically retrieved (**T**$_3$, orange). Fourth, the user analyzes variations of temporal progressions in the exploration phase from t_0 to t_1 (**T**$_4$). Finally, the user correlates the values of time series at time point t_1 (**T**$_5$) with external attributes.*

  - **T**$_3$: identify peers with most similar behavior before t_0
- **Exploration Phase**: explore variances within neighbors after t_0
  - **T**$_4$: explore commonalities and differences of peers after t_0
  - **T**$_5$: correlate post-progression of peers with external factors that may explain variations of peers after t_0

3. Visual Analytics Approach

We present LFPeers, a visual analytics system that adopts the conceptual workflow (Figure 2) in two dedicated interfaces, for the search phase (Figure 1) and the exploration phase (Figure 3).

3.1. Search Phase

The search interface provides an overview of the data objects in a table view (bottom left), attribute views showing temporal value distributions (on the left), a world map (top right), a scatterplot showing dimensionality-reduced countries (lower right), as well as controls for the definition of object similarity.

Visualization design The top left view allows users to browse through different groups of attributes and analyze temporal details of attributes. Details of temporal attributes are always represented with bundled line charts [JME10, BHR*]. As with the conceptual workflow in Figure 2, these bundled line charts follow a common color coding. The selected country (**T**$_1$) is drawn with a gradient blue-green line always on top. The gradient coloring mirrors the y-axis position and helps to distinguish the trends of the stacked sparklines within the dense layout. In LFPeers, the orange color is reserved for identified peers (**T**$_3$), here to support superimposed visual comparisons [GAW*11] with the selected country. All other countries are drawn in grey to provide the global context. For the temporal analysis of ordinal attributes, we modify the y-axis and use an ordinal scale, assigning each category to a certain bandwidth in the notion of a heatmap. The selected country is drawn as a line in the middle of these bands (see bottom of Figure 3). The *table* (lower left), shows the query object and nearest neighbors, in combination with preferred temporal and static attributes. Temporal attributes are displayed as sparklines, static attributes as color-coded rectangles. The table can be sorted by static attributes and by object similarity. This approach can be used as an entry point to identify countries with distinctive developments or trends.

Object selection The first action across all views is the selection of a single country of interest that is used as the query object (**T**$_1$). A combo box allows the selection by name (see Figure 1 top left), but users can also map attributes onto the world map, and select interesting countries from there.

Temporal querying LFPeers enables users to define an attribute and time range for the similarity calculation. This includes the meaningful definition of the endpoint t_0 of this time range, as t_0 will separate the search phase from the subsequent exploration phase (**T**$_2$). Every chart with a temporal x-axis can be used for the definition of the time range with a brushing interaction, inspired by the timebox widgets approach [HS04]. Time ranges are shown with a grey background across all temporal charts (see Figure 1 left).

Interactive similarity concepts Following the definition of a query attribute and time range, LFPeers computes feature vectors for all respective time series, and sorts the table by the similarity to the query object. LFPeers supports selecting between different time series descriptors to transform the data into numerical feature vectors, ranging from equidistant temporal sampling to statistical measures including min, max, or trend values. The set of descriptors is extendable, e.g., by PAA [KCPM01] and PIP [ZJGK10] algorithms. Users can also select between distance metrics: dynamic time warping [SC07] if similar structures shall be considered with a certain tolerance in temporal distortion, or Euclidean distance if the exact shape and temporal occurrences shall be emphasized [Ber15].

Identification of peers A natural way of identifying the peer group (**T**$_3$) is using the ordered list of countries in the table, allowing their detailed analysis and comparison. Conceptually, a peer group is empty at start and can be extended by adding countries of interest, by using the checkboxes for every row element. A straightforward approach may, e.g., be to add the first k neighbor objects, however, users also have the flexibility for more individual groupings. Candidate objects can also be added in the scatter plot using brushing interaction or by clicking on the map. In turn, both views can also be used to filter and fine-tune the peer group, e.g., to restrict the group to parts of the world or to remove outliers.

Multi-criteria peer selection To further support forming peer groups based on multiple criteria, LFPeers supports the definition

Figure 3: *An excerpt of the exploration phase interface, depicting the correlation between two time series that show noticeable fan-out behavior after the similarity selection range (marked by the orange rectangle below the x-axis). The orange dots depict the correlation of the two time series at multiple time points. The scatter plot presents a closer view at the data at the last time point.*

of similarity concepts for up to four attributes in parallel. Each concept can be based on entirely different definitions of time series similarity, leading to different orderings of countries. The final peer group is formed by the intersection of all user-defined country subsets. These parallel similarity concepts are visualized in the similarity selection view (Figure 1, center left) where the *Stringency Index* similarity concept is currently active.

Auxiliary views The components on the right provide additional structural and geographical context. The *scatter plot* depicts a t-SNE [VdMH08] projection of the feature vectors as one representative of the class of dimensionality reduction methods [EMK*19]. A crosshair marks the selected country and allows a more fine-grained assessment of neighborship relations. The point colors are assigned by their position in the 2D projection space [BSM*15] using a cube-diagonal RGB 2D colormap [Ves99]. To provide visual linkage to the table, an equally colored circle icon is added to every row. Given that the table is sorted by similarity, one would expect a range of similar colors across top-ranked objects. In turn, strong color variations indicate errors in the sorting process and a call for adjustments. The *world map* is also linked using the same coloring scheme, so that (data-wise) similar countries can be compared to their geographical locations. As an alternative, a binary color-coding can be used, highlighting the peer group across the globe.

3.2. Exploration Phase

An overview of all temporal attributes with bundled linecharts allows the exploration of peer development in later phases. Users can browse through a grid-based chart assembly and identify noticeable behaviors such as widespread fan-outs, or a divergence into sub-groups ($\mathbf{T}_4$). Once an interesting pattern is observed, users can select the attribute of interest, which is subsequently visualized at the top right and used as target attribute for correlation analysis. With the definition of a time point t_1 in one of the charts with a simple click interaction, LFPeers automatically computes the correlations between the target attribute and all other temporal and static attributes, by using the attribute values for all peers at time point t_1 ($\mathbf{T}_5$). A list interface on the right highlights the most interesting attributes by the order of their correlation strength. For clicked attributes of interest, LFPeers automatically shows a pop-over with details of the explaining attribute, including the distribution of the temporal attribute and correlation details shown in a scatterplot. Finally, the strength of correlations for alternative time points is shown at the bottom using a point size encoding. A time series that consistently displays very similar behavior might not be a candidate to explain e.g. a divergence into sub-groups, but might be affected by the same cause. A series that has shown less correlation up until t_0, but higher correlation after this point might be of more interest.

4. Usage Scenario

We demonstrate the applicability of LFPeers with a usage scenario across all five tasks using the Covid-19 case. An analyst seeks to identify distinguishing factors that drove the observations of Covid-19 cases in the US after the summer of 2020. The analyst selects the US as the country of interest ($\mathbf{T}_1$) and begins the analysis with brushing a time range ($\mathbf{T}_2$) until summer for the *total confirmed cases (relative)* (Figure 1, center left). Using the first similarity ordering, the analyst adds multiple countries to the peer group, as many countries in the grey time range appear very similar ($\mathbf{T}_3$). The world map displays his selection, and he deselects several countries to focus on Western and Asian regions. Looking at the set of peers (highlighted in orange), he sees increasingly varying behaviors in confirmed cases after the summer, with the US curve noticeably above almost all others. The analyst restricts the set of peers via two additional constraints *Stringency Index* and *Transit Stations*. The charts for those criteria show much more variability, so the analyst carefully browses through the sorted table to decide, which curve patterns to add to the set of peers. The projection hints towards several clusters; by brushing, he takes those as initial selection and continues to refine the selection, ending up with only 15 countries matching his interest. In the exploration page, he scans the charts for indications if the peers form patterns in government actions and external factors ($\mathbf{T}_4$). He quickly realizes that his peer group is too small for meaningful analysis, so he reiterates on his peer groups and continues with a larger group. He then probes autumn and winter time points in the chart of confirmed cases, triggering the system to compute and rank correlations ($\mathbf{T}_5$). He finally drills down into the detailed view (Figure 3), to investigate temporal relationships.

5. Discussion and Conclusion

We have presented a visual analytics approach for the cause-effet analysis in time-oriented data, using Covid-19 as the driving use case. At a glance, our approach supports users in the search for a group of similar data objects in an early phase of the time series, as well as in the exploration of variations within this group as occurring in later phases of the time series data. In addition, users have a means to explain these variations using correlation analysis. We believe that this type of approach complements the rich body of Covid-19 dashboards with a stringent focus on visual analytics capability and a clearly defined and yet unsupported analysis task.

Despite the relevance of finding solutions for the ongoing pandemic, our conceptualization already aims for a domain-agnostic and technique-driven agenda. With $\mathbf{T}_3$, we draw an explicit connection to searching time series by similarity, which comes with many degrees of freedom that have to be considered. We already allow users to steer the means to determine similarity interactively. Along these lines, there is ongoing work to extend LFPeers' toolkit, and apply the approach on other data to validate and reflect on the generalizability of the approach.

References

[AA06] ANDRIENKO N. V., ANDRIENKO G. L.: *Exploratory analysis of spatial and temporal data - a systematic approach*. Springer, 2006. doi:10.1007/3-540-31190-4. 2

[AGJ*20] AFZAL S., GHANI S., JENKINS-SMITH H. C., EBERT D. S., HADWIGER M., HOTEIT I.: A visual analytics based decision making environment for covid-19 modeling and visualization. In *IEEE Visualization Conference (VIS)* (2020), pp. 86–90. doi:10.1109/VIS47514.2020.00024. 2

[AKMR16] ABDELHAMID S. E., KUHLMAN C. J., MARATHE M. V., RAVI S. S.: Interactive exploration and understanding of contagion dynamics in networked populations. In *Conference on Behavioral, Economic and Socio-cultural Computing (BESC)* (2016), pp. 1–6. doi:10.1109/BESC.2016.7804480. 2

[AMST11] AIGNER W., MIKSCH S., SCHUMANN H., TOMINSKI C.: *Visualization of Time-Oriented Data*. Human-Computer Interaction Series. Springer, 2011. doi:10.1007/978-0-85729-079-3. 2

[BAP*05] BUONO P., ARIS A., PLAISANT C., KHELLA A., SHNEIDERMAN B.: Interactive pattern search in time series. In *Visualization and Data Analysis (VDA)* (2005), Erbacher R. F., Roberts J. C., Gröhn M. T., Börner K., (Eds.), vol. 5669, SPIE, pp. 175–186. doi:10.1117/12.587537. 2

[Ber15] BERNARD J.: Exploratory search in time-oriented primary data, December 2015. URL: http://tuprints.ulb.tu-darmstadt.de/5173/. 2, 3

[BHR*] BERNARD J., HUTTER M., REINEMUTH H., PFEIFER H., BORS C., KOHLHAMMER J.: Visual-interactive preprocessing of multivariate time series data. *Computer Graphics Forum (CGF) 38*, 3, 401–412. doi:10.1111/cgf.13698. 3

[BSH*16] BACH B., SHI C., HEULOT N., MADHYASTHA T., GRABOWSKI T., DRAGICEVIC P.: Time curves: Folding time to visualize patterns of temporal evolution in data. *Transactions on Visualization and Computer Graphics 22*, 1 (2016), 559–568. doi:10.1109/TVCG.2015.2467851. 2

[BSM*15] BERNARD J., STEIGER M., MITTELSTÄDT S., THUM S., KEIM D., KOHLHAMMER J.: A survey and task-based quality assessment of static 2d colormaps. vol. 9397, SPIE Press, pp. 93970M–93970M–16. doi:10.1117/12.2079841. 4

[BSW*14] BERNARD J., STEIGER M., WIDMER S., LÜCKE-TIEKE H., MAY T., KOHLHAMMER J.: Visual-interactive Exploration of Interesting Multivariate Relations in Mixed Research Data Sets. *Computer Graphics Forum (CGF) 33*, 3 (2014), 291–300. doi:10.1111/cgf.12385. 2

[BWK*13] BERNARD J., WILHELM N., KRÜGER B., MAY T., SCHRECK T., KOHLHAMMER J.: Motionexplorer: Exploratory search in human motion capture data based on hierarchical aggregation. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 19*, 12 (Dec. 2013), 2257–2266. doi:10.1109/TVCG.2013.178. 2

[CMP*19] CIBULSKI L., MAY T., PREIM B., BERNARD J., KOHLHAMMER J.: Visualizing time series consistency for feature selection. In *Journal of WSCG* (2019), pp. 93–102. doi:10.24132/JWSCG.2019.27.2.2. 2

[Com20] COMBA J. L.: Data visualization for the understanding of covid-19. *Computing in Science & Engineering 22*, 6 (2020), 81–86. 1

[DDG20] DONG E., DU H., GARDNER L.: An interactive web-based dashboard to track covid-19 in real time. *The Lancet infectious diseases 20*, 5 (2020), 533–534. 1

[EMK*19] ESPADOTO M., MARTINS R. M., KERREN A., HIRATA N. S., TELEA A. C.: Towards a quantitative survey of dimension reduction techniques. *Transactions on Visualization and Computer Graphics* (2019). 4

[GAW*11] GLEICHER M., ALBERS D., WALKER R., JUSUFI I., HANSEN C. D., ROBERTS J. C.: Visual comparison for information visualization. *Inf. Vis. 10*, 4 (2011), 289–309. URL: https://doi.org/10.1177/1473871611416549, doi:10.1177/1473871611416549. 3

[HS04] HOCHHEISER H., SHNEIDERMAN B.: Dynamic query tools for time series data sets: timebox widgets for interactive exploration. *Information Visualization 3*, 1 (2004), 1–18. 3

[JME10] JAVED W., MCDONNEL B., ELMQVIST N.: Graphical perception of multiple time series. *IEEE Transactions on Visualization and Computer Graphics 16*, 6 (2010), 927–934. doi:10.1109/TVCG.2010.162. 3

[KCPM01] KEOGH E. J., CHAKRABARTI K., PAZZANI M. J., MEHROTRA S.: Dimensionality reduction for fast similarity search in large time series databases. *Knowl. Inf. Syst. 3*, 3 (2001), 263–286. URL: https://doi.org/10.1007/PL00011669, doi:10.1007/PL00011669. 3

[KH13] KEHRER J., HAUSER H.: Visualization and visual analysis of multifaceted scientific data: A survey. *IEEE Trans. Vis. Comput. Graph. 19*, 3 (2013), 495–513. doi:10.1109/TVCG.2012.110. 2

[LSC*20] LEITE R. A., SCHETINGER V., CENEDA D., HENZ B., MIKSCH S.: COVIs: Supporting temporal visual analysis of covid-19 events usable in data-driven journalism. In *IEEE Visualization Conference (VIS)* (2020), pp. 56–60. doi:10.1109/VIS47514.2020.00018. 2

[MA14] MIKSCH S., AIGNER W.: A matter of time: Applying a data–users–tasks design triangle to visual analytics of time-oriented data. *Computers & Graphics 38* (2014), 286–290. doi:https://doi.org/10.1016/j.cag.2013.11.002. 2

[PBK*19] POLONSKY J. A., BAIDJOE A., KAMVAR Z. N., CORI A., DURSKI K., EDMUNDS W. J., EGGO R. M., FUNK S., KAISER L., KEATING P., DE WAROUX O. L. P., MARKS M., MORAGA P., MORGAN O., NOUVELLET P., RATNAYAKE R., ROBERTS C. H., WHITWORTH J., JOMBART T.: Outbreak analytics: a developing data science for informing the response to emerging pathogens. *Philosophical Transactions of the Royal Society B: Biological Sciences 374*, 1776 (2019), 20180276. doi:10.1098/rstb.2018.0276. 2

[SC07] SALVADOR S., CHAN P.: Toward accurate dynamic time warping in linear time and space. *Intelligent Data Analysis 11*, 5 (2007), 561–580. 3

[Tim20] TIMES F.: Coronavirus tracked: the latest figures as countries fight to contain the pandemic. *Financial Times 12* (2020). 1

[VdMH08] VAN DER MAATEN L., HINTON G.: Visualizing data using t-sne. *Journal of machine learning research 9*, 11 (2008). 4

[Ves99] VESANTO J.: SOM-based data visualization methods. *Intelligent data analysis 3*, 2 (1999), 111–126. 4

[W*20] WAHLTINEZ O., ET AL.: COVID-19 Open-Data: curating a fine-grained, global-scale data repository for SARS-CoV-2. Work in progress. URL: https://goo.gle/covid-19-open-data. 2

[WGGP*11] WONGSUPHASAWAT K., GUERRA GÓMEZ J. A., PLAISANT C., WANG T. D., TAIEB-MAIMON M., SHNEIDERMAN B.: Lifeflow: Visualizing an overview of event sequences. In *SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2011), Association for Computing Machinery, p. 1747–1756. doi:10.1145/1978942.1979196. 2

[WR09] WHITE R. W., ROTH R. A.: Exploratory search: Beyond the query-response paradigm. *Synthesis lectures on information concepts, retrieval, and services 1*, 1 (2009), 1–98. 2

[YDH*17] YAÑEZ A., DUGGAN J., HAYES C., JILANI M., CONNOLLY M.: Pandemcap: Decision support tool for epidemic management. In *IEEE Workshop on Visual Analytics in Healthcare (VAHC)* (2017), pp. 24–30. doi:10.1109/VAHC.2017.8387497. 2

[ZJGK10] ZIEGLER H., JENNY M., GRUSE T., KEIM D. A.: Visual market sector analysis for financial time series data. In *IEEE Symposium on Visual Analytics Science and Technology* (2010), pp. 83–90. doi:10.1109/VAST.2010.5652530. 2, 3