

Visual Analytics of Contact Tracing Policy Simulations During an Emergency Response

M.Sondag^{1,5} , C. Tur kay² , K. Xu³ , L. Matthews⁴ , S. Mohr⁴  and D. Archambault¹ 

¹ Swansea University, United Kingdom

² University of Warwick, United Kingdom

³ Middlesex University, United Kingdom

⁴ University of Glasgow, United Kingdom

⁵ University of Cologne, Germany

Abstract

Epidemiologists use individual-based models to (a) simulate disease spread over dynamic contact networks and (b) to investigate strategies to control the outbreak. These model simulations generate complex ‘infection maps’ of time-varying transmission trees and patterns of spread. Conventional statistical analysis of outputs offers only limited interpretation. This paper presents a novel visual analytics approach for the inspection of infection maps along with their associated metadata, developed collaboratively over 16 months in an evolving emergency response situation. We introduce the concept of representative trees that summarize the many components of a time-varying infection map while preserving the epidemiological characteristics of each individual transmission tree. We also present interactive visualization techniques for the quick assessment of different control policies. Through a series of case studies and a qualitative evaluation by epidemiologists, we demonstrate how our visualizations can help improve the development of epidemiological models and help interpret complex transmission patterns.

CCS Concepts

• **Applied computing** → **Health informatics**; • **Human-centered computing** → **Visualization design and evaluation methods**; **Visual analytics**;

1. Introduction

In March 2020, COVID-19 became a pandemic which prompted the introduction of a variety of different and often intrusive policies by various governments to curb the spread of the disease. During this time, the scientific community was called to action to understand the effectiveness of these policies before and while they are being applied and to model future policy scenarios. One type of model used by epidemiologists is an individual-based model where the disease spreads through inter-person contacts. The contacts are modelled as a temporal network [HS12] – a network where each edge has its own timestamp when the contact occurred. If an infectious individual comes into contact with a susceptible individual, the disease can be transmitted. Policies can be modelled on top of these networks to evaluate their effectiveness. The result is a forest of disconnected trees, or an *infection map*. Each individual tree in these maps models the varying dynamics and topology of infections in a “local” network. In addition, the nodes and edges carry associated metadata, such as the location that the transmission took place (e.g., workplace, restaurant, school, etc.), the disease status of the infector at the time of infection (e.g., pre-symptomatic or symptomatic) or the demographics (e.g., age of the individuals). Even though this metadata is valuable for both validating models

and comparing multiple policy scenarios, the interpretation of these model results often resorts to statistical analysis that aggregate, or sometimes even discard, the rich metadata within these trees. Instead, they focus solely on the eventual aggregate indicators, such as the number of deaths or hospitalizations. This reliance on aggregate statistics is primarily due to a lack of effective visualization approaches that can support the epidemiological modelling workflow through an in-depth analysis of the infection maps. This paper responds to this gap in knowledge.

Our visualization team began its interactions with the epidemiologists in May 2020, to create visualization solutions to help with the setting of model parameters and making sense of the output of model runs. The primary task here was to create visualizations to help understand the model outputs and the effect of model parameters and policies. Our 16 months engagement has led to a series of prototypes, including the one presented in this paper, and also a way of interacting with end users in emergency response situations. As part of our contribution, we document our interactions with our end users in this emergency scenario.

The primary contribution of this paper is a novel technique that visualizes infection maps of disease spread using the concept of

representative trees. These representative trees cluster local outbreak patterns based on their similarity in disease spread and allow for an aggregated overview while preserving epidemic structures for large infection maps. The technique was developed in collaboration with epidemiologists interested in understanding the effect of parameters and policies for their individual-based simulator of COVID-19. We evaluate and demonstrate the effectiveness of our approach through a series of use cases and a qualitative evaluation with epidemiologists where tasks supportive of their modelling workflow are conducted. Through these evaluation activities, we observe that our visualization approach facilitates richer and effective interactions with model outputs that enables epidemiologists to better understand the results of the models, identify issues with the models, explore the impact of parameters on disease outcomes and policy effectiveness, and allow for comparison and identification of effective policies that reduce disease transmission.

2. Background and Related Work

In this section, we will briefly introduce epidemiological modelling and related visualization techniques.

2.1. Disease Modelling and Individual-Based Models

Contact tracing and testing have a direct impact on epidemic control (e.g. through isolation of contacts of symptomatic individuals) and improve our understanding of control measures and wider disease transmission in general [JPH*21]. There are a number of models that explicitly represent how contact tracing together with other non-pharmaceutical interventions affect the dynamics of COVID-19 from various perspectives: the effectiveness of digital contact-tracing [FWK*20, HPN*20], the impact of quarantine length and test-and-release strategies [QCH*21], heterogeneous contact networks as drivers of the epidemic [FWK*20], or the required speed for effective contact-tracing [He20, Ke20, KRB*20].

Typically, aggregate statistics are used to quantify the epidemiological transmission potential in specific settings. These aggregate statistics can be useful for assessing the overall societal burden of contact tracing, but do not allow for the investigation of the impact of complex contact-tracing policies on individual-level transmissions. In our visual analytics approach, we pair an individual-based model with a visual analytics system. Similar visual analytics systems have been developed for individual- and agent-based models, but usually they are aimed at models with hundreds of nodes [MFM*18, GBA20] and do not help users understand the network structures at an aggregated level [CAD*14], or have explicit spatial components in the data [BWMM15, GPLR21, WXL*21].

Model to be visualized. The COVID-19 Contact-Tracing Model (developed as part of the COVID-19 RAMP response) visualized in the paper is an individual-based stochastic network model used to explore flexible contact tracing and testing policies [git] originally developed from Mohr *et al.* [MDC*18]. The model for spreading the disease over the contacts is a modified SEIR model with asymptomatic and pre-symptomatic transmission. It operates on the underlying network of dynamic contact data, a temporal network without spatial components, stratified by setting (home, work, school, etc.) to spread the disease. Nodes in this network represent

individuals classified by age; edges occur at a specific time and are represented by weighted contacts where the weight represents the duration of a contact. Parameters used in the model were derived from literature [CBGM*21, BMC*20, GCC*20, MCH*20]. The result of the simulation is an infection map: a forest of rooted directed *transmission trees*. Each transmission tree has a root node representing an *index case* that was exposed to the disease at the start of the simulation. A directed edge $e = (u, v, t)$ in the tree indicates that node u has transmitted the disease to node v at time t .

2.2. Visualization for Public Health and Pandemics

Visual analytics systems have been used in the public health domain, for example through health records [HHO*16, GXZ*18, RWA*13, MLL*13] and epidemiological data [PL20]. For infectious disease control in particular, visual analytics approaches have been used for simulated disease data from spatio-temporal perspectives [BWMM15], used as decision support tools for pandemic management [MLR*11, YDH*17] and used to trace back disease outbreaks to their origin in hospitals [BPW*21, MPW*20].

For COVID-19 specifically, a number of visual analytics systems for situational awareness and policy decisions have been proposed [MHW*21, DHA*20, LSC*20], including from a geospatial perspective for simulations [AGJS*20]. In terms of visual analytics systems to support contact tracing visualizations, there has been some work that has integrated link prediction with visualizations to visualize potential clusters of COVID-19 contacts [ASGK21]. Although there has been significant work on the dynamics of a pandemic from a variety of perspectives at a high level, this work focuses on visual analytics to support the dynamic relationships in an individual based simulation supporting contact tracing. An understanding of these simulations can help with understanding and developing new contact tracing policies.

2.3. Visualizing Graphs of Many Components

Methods have been created to summarize collections of graphs [LSDK18]. Techniques have been explored and evaluated for summarizing graphs and for labeled trees where one-to-one matching [Arc09, APP10, MGT*03] or limited matching via node label [KZA10] is available. Koop *et al.* [KFS13] superimposes a number of labeled graphs to construct a visual summary. The labels here provide a one-to-one correspondence for all graphs in the set, meaning that they can be overlaid onto a supergraph containing all nodes and edges from the set. However, for our specific problem no such labels are available and hence graphs cannot be matched directly in this way. Additionally, making a supergraph (or in our case, a supertree) has the disadvantage that it contains a larger number of nodes than the trees from which it is composed, and the number of nodes is an important indicator of epidemic impact.

Graphlets is a similar, but slightly different method compared to our technique. Graphlets are small ($n < 5$) subgraphs of a larger graph through which the structure of larger graphs can be characterized [RPS*21]. They have been used to gauge structural similarity in egocentric networks for visualization [HACH12]. While these graphlets could be used similarly to summarize transmission trees, they lend themselves less easily to visualizing the metadata of the

nodes and viewing the structure of the trees themselves. Moreover, being limited to small subgraphs is problematic, as the most interesting trees are those that have many transmissions.

Methods to visualize highly disconnected graphs by clustering the components on structural similarity have also been devised [SBTK08, vLGS09]. Statistics or features for each component are recorded and self-organising maps (SOMs) are used to detect clusters of components with similar attributes. Although we solve a similar problem in this paper and suggest representatives for each cluster of components, our components are dynamic instead of static, and we use domain-specific knowledge for clustering.

3. Emergency Response: Early Days to Prototypes

The design of visualizations in this paper is the result of a collaboration of an interdisciplinary team of domain experts in epidemiology and visualization. Over a 16-month period, we conducted bi-weekly meetings during which we discussed the models, their parameters, the data emerging from the models, as well as key policy questions in relation to the modelling results. At the beginning of our collaboration in May 2020, we did not have the opportunity to select our collaborators or the problem: our contact tracing team was formed through a national-level organization, RAMPVIS through the Royal Society, for the COVID-19 response; all we had was our enthusiasm to help and our respective expertise. Also, the expectation was to respond quickly to the modelling needs by developing visualization prototypes in an environment where the requirements and priorities were constantly changing (our definition of emergency response and that of RAMPVIS [CARA*21]). A carefully designed visualization would have taken too long and would be less useful in an evolving pandemic situation. The epidemiologists had limited experience with visualizing graph data; one mentioned that they had not known of visualization as a discipline before, and the visualization experts had limited knowledge in epidemiology. In normal circumstances, the appropriate way to approach the design process would be to adopt the design study methodology (DSM) [SMM12] and rigorously follow its nine stages. However, due to the challenges outlined above, we significantly shortened its initial stages.

Our emergency response design study methodology did not include any of the stages before the *discover* stage [SMM12]. The reason for the exclusion of these stages is due to the fact our collaborator and problem were selected for us in an emergency situation. This leaves us prone to several pitfalls of the design study methodology, in particular: *unsuitable problem due to potential for automation* and *collaboration with the wrong people*. However, these risks were necessary given the situation. These modifications are similar to those in the lite methodology [SMR*20], the *choose a collaborator* stage does not exist, as our collaborators could not be selected due to the nature of the response. Secondly, the *abstract* phase has been heavily modified.

Our methodology begins at a heavily modified *discover* stage (known as the *abstract* phase in the lite methodology) where we focus on problem analysis and task abstraction. We began this stage by meeting with our collaborators and understanding the data produced by the contact tracing simulation. As our collaborators had no experience working with visualization researchers, we

performed rapid prototyping with existing visualization tools that were at our disposal: GMap [GHK10] to demonstrate areas in the data where the infection location was the same, temporal dynamic graph animations of the data drawn in the space-time cube [SAK20, SAK17, AMA21], time-to-colour encodings [BDA*17] of the data in Jupyter notebooks, and using tools such as Gephi [BHJ09] (see supplementary material for prototypes). These pre-prototypes provided an initial grounding in the visualization field and allowed us to begin work on visualizations that fit their needs, finding out what worked and what did not work, and most importantly, as mediums to build a healthy dialogue within the team.

From our discussions and inspection of the related literature, we observed that epidemiologists rarely use visualizations more complex than line charts to investigate their data. Moreover, intermediate artifacts such as how the disease spreads throughout their model were underutilized. Thus, this initial application of existing visualization tools to their simulation data was crucial for buy-in and turned out to be very successful. Our collaborators were quite interested in these results and expressed their appreciation publicly during RAMP consortium meetings. To our surprise, even these initial design explorations using existing tools were helpful in leading to useful findings in relation to the simulation data, and several bugs and inconsistencies in the simulation were discovered. For example, in early runs of the simulator the number of random infections (infections from outside sources) was set to too high a value, giving unexpected results, confirming that visualization techniques can help with debugging and analyzing the results. We also found that the *cast* phase [SMM12] had significantly more informal roles.

Our discovery phase found that the graph structure of the infection map was atypical (see supplementary material). Instead of the graph consisting of few components, it consists of a large forest of trees with many small trees and a few larger trees. The epidemiologists are mostly interested in the larger trees, as these show patterns where the disease spreads, providing a basis to tailor policies for preventing spread at a larger scale. However, the large number of small trees should not be simply discarded. While a single small tree does not hold interesting information, many small trees combined do as they for example indicate that the disease has low reach in many areas of the infection map.

Subsequently, we entered more traditional core phases of agile, inward-facing validation. We gradually moved away from existing prototypes and moved towards custom prototypes as the tasks of our end users became clear.

4. Visualization Problem and Tasks

During the explorative phase of the project outlined above, we followed Roger et al.'s [RPH*20] recommendations for rigorous and rich note-taking, and made use of regularly revised project diaries that mix visualizations and narrative. These were shared on an open repository and used as mediums for exchange during the meetings. Through these interactions, we consolidated our understanding and learnings in the form of a number of requirements that a useful visualization for our epidemiologist collaborators should adhere to.

The input for our problem is a stochastic simulator developed by

the epidemiologists and their collaborators that generates an infection map, including COVID-prevention policies that can be modelled using the simulator. We identified three main requirements for the visualization. First of all, *the visualization should be able to give a concise overview of the infection map*, which often contains tens of thousands of nodes. Secondly, *we should be able to visually explore the metadata for each node*. Thirdly, *we should be able to show the impact of different policies on the same infection map*. By using the same infection map as a base instead of comparing multiple runs of the model, we aim to investigate different policies quicker and allow for a direct comparison of potential policy impact that without being biased by the inherent randomness in the simulation. Aside from these requirements, we consider three higher level tasks that our visualization should support, aimed at the middle phases ("Model validation", "Comprehension of causes and effects within the model" and "Perception and detection of patterns in a simulation for later formal quantification") of an agent-based modelling project as Dorin et. al's [DG14] identifies:

- T1 Observe and confirm whether the model is behaving as expected. Model parameters for new diseases are usually adapted rapidly to handle a new outbreak, and this can involve a lot of potentially error-prone parameter tuning. The epidemiologist has a clear mental model of the simulator and some expectation of the results. The visualization of the behavior of the model should match their mental model and expectation. If not, it means that either the simulation requires further parameter tuning (e.g., there is some error in the current setting) or there are some unexpected behaviors that require further investigation.
- T2 Explore how the disease spreads through the network. Providing insight into the dynamics of the transmission of the disease helps epidemiologists in coming up and developing control policies to prevent the spread of disease. By giving a more informative insight via a visualization, this is both easier to communicate to policymakers and allows for more detailed understanding of the behavior of the disease and the impact of a specific policy.
- T3 Compare the effect of policy settings. The policies considered by the epidemiologists all have multiple parameters such as when a person isolates, how long a person isolates, how many contacts of the person are warned given a positive test, etc. The epidemiologists should be able to quickly compare the effect of different settings to understand the role and impact of the settings.

5. Visual Analytics of Infection Maps

In order to analyze many transmission trees, we designed a visual analytics system that clusters trees based on epidemic similarity. We calculate representative trees for each cluster, and visualize the effect of policies for these clusters on these trees.

5.1. Representative Trees

There are many disconnected components and nodes, and hence it is not feasible to display all the data using a standard node-link diagram. Therefore, we cluster the trees in such a way that the aggregated structure characterizes the set of trees it represents. We achieve this via the concept of a *representative tree*. A representative tree represents a number of different trees with similar epidemic structure. These representative trees can be viewed as the




Figure 1: Two transmission trees with the time of infection inscribed within the nodes. The R_t values of the trees are (a) $[\frac{2}{1}, 0, 0, 0, \frac{0}{1}, 0, 0, 0, \frac{1}{1}, \frac{0}{1}]$ and (b) $[\frac{1}{1}, 0, \frac{2}{1}, 0, 0, 0, 0, 0, 0, \frac{0}{2}]$. Smoothing the R_t values over 7 timesteps gives (a) $[\frac{2}{7}, \frac{0}{7}, \frac{1}{7}, \frac{1}{7}]$ and (b) $[\frac{3}{7}, \frac{2}{7}, \frac{2}{7}, \frac{0}{7}]$ for a smoothed R_t distance of $\frac{5}{7}$.

center of a cluster, where the trees within the cluster have a small "epidemic distance" to this center. To compute these representative trees, we first define a distance measure to measure the similarity between two trees. A commonly used tree distance measure is the tree-edit distance that measures how many insertions and deletions are needed to turn one tree into another under some constraints. However, trees with low tree edit-distance can be quite different in how a disease spreads through them over time. Moreover, the epidemiologists in our team found the tree-edit distance difficult to grasp during testing. Therefore, we introduce the R_t distance measure, which measures how similar the disease develops in two trees.

We base the distance between two trees on R_t (more commonly known as R as used in the media), the effective reproduction number [NC09]. This parameter is often used by epidemiologists to capture how a disease spreads over time. As we have access to whom infected whom in the infection map, we can calculate R_t for each time t and each transmission tree T as follows [STM*22]:

$$R_t = \begin{cases} 0, & \text{if } |X(t)| = 0 \\ \frac{|Y(t)|}{|X(t)|}, & \text{otherwise} \end{cases}$$

where $X(t)$ is the set of all nodes that are newly exposed at time t , $|X(t)|$ is the amount of nodes newly exposed at time t , and $|Y(t)|$ is the amount of children of $X(t)$. Note that while R_t can be calculated as a total over all trees as well, we only calculate it for each tree separately in this paper. To decrease the effect of noise when calculating distances between trees, we average the value of R_t over a time window. In our implementation, we smoothen over 7 timesteps. Using these smoothed averages, we calculate the R_t distance between two trees by summing the absolute difference between the smoothed R_t values of both trees for each timestep t . The smaller this distance, the more similar the trees are from an epidemiological viewpoint. This is illustrated in Figure 1.

With this distance measure, we can calculate the representative trees and their associated clusters. We do not set a fixed level of aggregation, as different data and different tasks require different levels of aggregation. Instead, we allow the user to specify the maximum R_t distance within the trees of a cluster. With a small distance, few trees will be aggregated and many representative trees will be required, allowing for a detailed view. By increasing the distance, the user can reduce the amount of representative trees on-screen and aggregate them to gain a higher level view of the data.

As the user changes the distance parameter, representative trees can appear and disappear in the visualization. To prevent trees from flickering in and out of the visualization, representative trees should




Figure 2: Graph of R_t distances between representative trees (left). Graphs G_d for $d = \{0, 1, 2, 3\}$ with nodes in dominating sets D_d colored black, and edges not in G_d grayed out (right).

be stable with regard to the changing distance parameter: when increasing the distance parameter, trees should only ever appear and never disappear. This is achieved by finding a base set of representative trees and removing representative trees from this set as the distance increases as long as all trees remain represented (this will be explained in more details shortly). For each distance d , we represent our forest of transmission trees as a graph G_d . Each tree x in the forest represents a node $n(x)$, and there is an edge in G_d between two nodes $(n(x), n(y))$ if node $n(x)$ has an R_t distance of less than d from node $n(y)$. Moreover, x and y should have the same number of nodes to prevent smaller trees from being clustered together, as our domain experts indicate that there is an epidemiological difference between trees with different amount of nodes.

To find our base set of representative trees, we calculate a dominating set D_0 for distance 0 on the graph G_0 . A dominating set D is a set of nodes such that every node not in D is adjacent to D . We say that a node v dominates a node u if u is adjacent to v and $v \in D$. To calculate our base set of representative trees, we select a representative tree y for each tree x , such that $n(x)$ is dominated by $n(y)$ and the R_t distance between x and y is minimal among all possible dominating nodes $n(y) \in D$. We then increase the distance value d by 1 and calculate D_1 . Initially, $D_1 = D_0$. We then iteratively go through each node $n(y) \in D_1$, removing $n(y)$ from D_1 if $D_1 \setminus y$ remains a dominating set of G_1 . After each removal, we reassign any nodes $n(x)$ that were dominated by $n(y)$ to another node $n(z)$ in D_1 that dominates $n(x)$. We incrementally increase the distance d by 1 and repeat this process until d equals the maximum R_t distance between two nodes. When this process completes, any tree x will be represented by a representative tree y for any distance d . Moreover, as we only delete nodes from the dominating set when increasing the distance, no new representative trees appear, and hence our representative trees are stable with regard to the distance parameter. An example of this process is shown in Figure 2.

Finally, we map the vertices of each tree x to its representative tree y by finding a mapping with minimal tree edit distance which preserves the structure of the tree as best as possible. As a result, each node in a representative tree represents a number of nodes from different transmission trees, whose metadata we visualize using a stacked bar chart in our visualization. The concept of representative trees and the visualization of them is shown in Figure 3.

Using these representative trees, we can explore how the disease spreads throughout the network by investigating the structure of the trees (T2). By visualizing the metadata, the domain experts




Figure 3: (a) Four transmission trees. The number in each node denotes the time of infection. The color of a node denotes which node in the representative tree it is mapped to. White nodes are not represented as they are deleted in the minimal tree edit distance mapping. (b) Representative tree of the transmission trees. (c) Visualizing the infection time of the represented nodes using a stacked bar chart. The time stamps for the middle represented nodes are 8, 2, 2, 7, thus the aggregated node has light blue in the bottom half (two '2's) and two dark sections on top (for the '7' and '8').

can verify that the model is working as expected and there is no unexpected behavior (T1). Thus, we propose that the visual analytics design above supports the first two tasks. We describe how to support the final task (T3) in the next section.

5.2. Displaying Policies

The remaining task that we set out to solve is to compare different policy settings (T3). Running the full simulation to determine the effects of changing a setting can change 'whom-infects-whom' between scenarios due to the stochastic nature of the simulation. Therefore, in order to visualize the differences in transmission between settings, we apply a policy on the infection map of a baseline scenario, i.e., without any policy, in a post-hoc manner.

This reflects a retrospective look at what could have happened if a policy had been applied on a particular real-world example. It does have the disadvantage that it does not give a completely accurate view of how effective a policy is on its own: while a transmission path from person u to person v might be prevented on the infection map due to the policy, a different transmission path to infect person v can exist while not being in this infection map. The choice of applying policies post-hoc is thus a trade-off between better understanding *how* policy settings influence the simulation, and getting the complete picture of the effect of a single policy.

To facilitate the comparison of two policies' settings, we adapt the visualization of the representative trees. Instead of showing a single stacked bar chart per node of the representative tree, we show two stacked bar charts, one for each policy setting, such that they can be compared. We encode the policy using a separate color (green) that is not shared with the color schemes for the attributes, which allows us to combine visualizing attributes with visualizing policies to better understand the impact of a setting. An example of a representative tree being used to compare two policy settings A (left stacked bar




Figure 4: Comparing two policy settings (green) and time of infection (blue).

chart) and B (right stacked bar chart) along with the time of infection is shown in Figure 4. The green bars in a representative node n indicate how many transmission events have been prevented by a policy setting for the nodes represented by n . In this example, policy setting A is preventing slightly more transmission events at the low levels in the tree, and many more transmission events at the higher levels in the tree compared to policy setting B. To demonstrate how we support all three tasks, we work through a number of use cases with these representative trees in Section 7.

6. Implementation

The implementation of our tool is split into two parts, both available online [imp]. In the first part, we preprocess the data using Java by calculating the representative trees. In the second part, we use JavaScript and D3.js [BOH11] to visualize the representative trees.

To calculate the representative trees for the first part of our implementation, we need approximate solutions to two NP-hard problems: calculating a minimal dominating set [Kar72], and calculating the mapping that minimizes the tree-edit distance [Bil05]. To calculate a dominating set, we use a simple greedy approximation algorithm: iteratively add the node which dominates the most non-dominated nodes to the dominating set. An approximate solution is sufficient here, as the dominating sets are necessarily not minimal due to our stability requirement when we increase the distance.

For the mapping of trees to their representative trees, we use the linear program as presented by Kondo et al. [KOIY14] to calculate an exact solution. While the problem is NP-hard, the linear program is fast enough in practice as most of our trees are relatively small. Thus, while inexact algorithms exist [AFH*13] that can be used for larger trees and datasets, they are not required for our use case. Using IBM CPLEX 20.1.0 as our solver and a standard computer with 16GB memory, we can map two trees with up to 24 nodes within a second, 46 nodes within a minute, and only starting from 74 nodes it starts becoming infeasible for our use case taking more than an hour to map two trees. As we have few large trees in the data, and we can preprocess the data, we can calculate most mappings exactly. For the largest pairs of trees (19 out of 5,000 mappings) the LP runs into the timeout limit we set of 1 hour, and we use the best solution the LP has found so far instead of an optimal solution. If more scalable approaches are required, other mapping algorithms for rooted unordered unlabeled trees could be also be used [Bil05].

As all policies and parameter settings that the epidemiologists were interested in (3 different durations of self-isolation, 4 different time periods for backward contact tracing, and 10 different percentages of app-uptake, for a total of 120 settings combinations) were known in advance, we preprocessed the calculations and stored the results to use in the second part of our implementation. However, calculating a single policy post-hoc is quick, and it would be fairly straightforward to calculate these interactively if required.

In the second part of our implementation, we visualize the representative trees using two stacked bar charts for each node, which requires around 15,000 svg elements. In our implementation, there is a small delay before the visualization gets updated. While this delay could be mitigated by further optimization, it was not deemed a hindrance from our expert user feedback.

6.1. Visual Design

Figure 5 shows the interface of our visual analytics approach. In area 1, we visualize the infection map using representative trees. The number of trees in the cluster is shown below each representative tree. Clicking on a representative tree reveals all the trees it represents (area 1a).

In area 2, we can select different visualization settings. In area 2a, the maximum R_t distance within a cluster is set. A scented widget [WHA07] is provided showing how many representative trees will remain on-screen for each distance. Adjusting the slider gives immediate feedback by fading out the representative trees that will disappear with the new setting. In area 2b, the base size of nodes can be adjusted with the layout of the representative trees adjusting as required. Area 2c specifies the node properties to be visualized on the left and right stacked bar charts of each node. In area 2d, policy settings and the parameter setting app-uptake (how many people have a COVID-tracking application) can be selected for visualization. When selecting a specific policy, nodes that are prevented from being infected in this infection map by these settings are given a green color. A "detailed" checkbox has been included after feedback from the domain experts that splits the prevention into people whose infector is isolating due to the policy, and people who no longer get infected due to the transmission chain being broken earlier in the chain. Area 2e has a button that launches the recalculation of the layout of the trees and the color schemes. It is not automatic as it takes a few seconds to recalculate the representative trees, which could confuse the user when the visualization is changing when they are still changing other settings.

Area 3 shows a global overview of the data and a color legend. In area 3c, the legend for the current selections in area 2 is shown (colors selected using colorbrewer [HB03]). Area 3b shows the distribution of node properties using the settings in area 2 through stacked bar charts. This distribution can either be from all nodes or all nodes from a specific level in the trees which can be selected via the two selectors in 3a. The user can decide to show absolute values or normalize them using a toggle button.

7. Use cases

To demonstrate how our proposed solution supports the underlying tasks (Section 3) and reveals relevant observations to an expert user (denoted as E in this section), we present three use cases, each of which addresses one of the tasks.

Use case 1: For the first task (T1: Observe and confirm whether the model is behaving as expected) user E will verify that the compartmental model overlaid on the temporal network is working correctly. For this, E first visualizes the infector state. Looking at the global overview (Figure 6a), E sees that the distribution of infectious states is as expected. Next, E looks at the representative trees and picks one of the larger trees (Figure 6a) to verify that the nodes go through all compartments, which is indeed the case. Moreover, E sees that the root node of the tree correctly goes through one branch of the compartmental model having infected nodes in the *Presymptomatic*, *Symptomatic* and *Severely Symptomatic* states, and none in the *Asymptomatic* state which would have indicated an error in the model. Finally, E verifies that there is no strange behavior in the




Figure 5: A screenshot from our tool visualizing simulation data. Different areas are highlighted using red rectangles and labeled. Area 1 shows the visualization of representative trees. Area 2 shows the interface for filtering and selecting attributes for the representative trees. Area 3 shows a high level overview of the distribution of the attributes and a color legend.

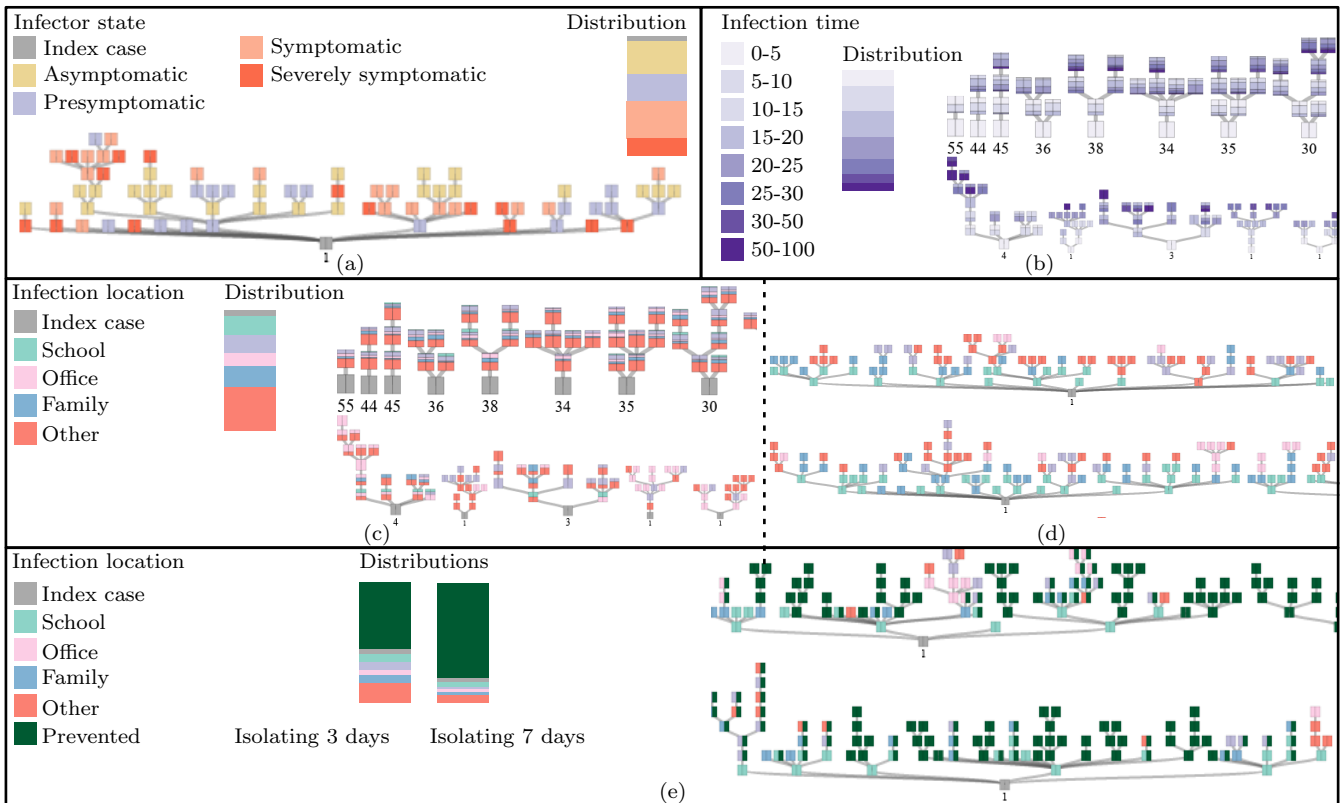


Figure 6: (a) Nodes correctly go through the states of the compartmental model. (b) Nodes have a consistent window in which they infect others. (c) School infections are underrepresented in smaller trees compared to the overall distribution. (d) School infections are the instigators of most of the larger trees. (e) Comparing policy settings by increasing isolation time from 3 days (left bar chart) to 7 days (right bar chart). Increasing isolation time has a large overall effect including on the larger trees, but some subtrees from the larger trees remain.

duration between infection times by visualizing these times. For example, an overly consistent coloring between siblings, or too large jumps in colors between a parent and a child node would indicate to *E* that there are errors or unknown insights in the underlying model.

Use case 2: We illustrate the second task (T2: Explore how the disease spreads through the network) by letting *E* explore the location of a transmission. When *E* looks at the overall distribution (Figure 6c), it is clear that school transmissions should have a decent presence in the representative trees. Yet, when *E* looks at the smaller trees (Figure 6c), there are barely any school transmissions. When *E* scrolls down and the trees become larger, more and more school transmissions (Figure 6d) become apparent, but these transmissions are generally at the start of the tree. This indicates to *E* that school transmissions have a larger impact than they appear to have from the distribution: Infections at schools often seem to result in large outbreaks, which spread further into the community, and thus result in more indirect infections than other places.

Use case 3: We illustrate the third task (T3: Compare the effect of policy setting) by letting *E* compare the effect of two policy settings on school transmissions: a policy setting where the time a person isolates is set to 3 days, and a policy setting where this time is 7 days instead. In Figure 6e *E* sees in the bar charts that this extra isolation time has a large impact overall. When *E* looks in more detail at the larger trees, the extra isolation time seems to have a large impact on preventing school transmissions for most of these, but for some subtrees it seems to have no effect at all. *E* hypothesizes that this is due to the root node of the subtree being asymptomatic, and this is indeed the case when *E* switches view to verify their hypotheses.

8. Qualitative Expert Evaluation

In order to understand and assess how epidemiologists use our tool in conducting the analytical tasks, as well as gather reflections on the analytical process facilitated by the tool and on the emergency design process, we report here on a qualitative evaluation conducted. Our evaluation approach can be considered as an instance of the "Evaluating Visual Data Analysis and Reasoning" scenario within Lam et al. [LBI*11] framework for visualization evaluation. We conducted the evaluation in two sessions with two epidemiology experts in each. The two participating experts in the first session (indicated as **E1** and **E2**) are also co-authors on this publication and participated in the co-design process, the other two experts (**E3** and **E4**) who joined the second session encountered the tool for the first time in the evaluation session.

Both evaluation sessions followed the same structure and the same script, although the second session deviated from the script as **E3** and **E4** were seeing the tool for the first time and had to ask more questions to get an understanding of the underlying model, the approach, and the tool. The evaluation session was split into two parts, first a think-aloud case study with a series of activities to conduct, followed by a semi-structured interview reflecting on the design, the design process (only in the first session), and the role of visualization in epidemiological modelling. One of the visualization researchers introduced the activities, asked prompting questions, responded to practical queries, and asked the interview

questions, while all visualization researchers took notes during both parts. Both parts were done online via Zoom. In the following, we report the results from both the sessions after a joint analysis.

8.1. Think-Aloud Case Study

For the first part of the evaluation, the experts conducted four activities that require them to make use of the tool while following a thinking-aloud process [AAP*05, Lew82]. Participants were specifically asked to talk through their analytical thinking and reasoning in relation to the actions they were taking. The session took place online over an hour with one of the experts sharing screen with the tool open, while the analysis was steered collaboratively by both experts. The experts were tasked to carry out the following four activities that are aligned with the analytical tasks in Section 3:

- A1:** *How would you use the tool to assess the model, i.e., is it doing what you expect it to do?* (aligned with **T1** and **T2**)
- A2:** *Are there any interesting patterns that you spot in the simulation about how the disease spreads? If so, why?* (**T1** and **T2**)
- A3:** *Would you be able to compare two policies that are of interest?* (**T2** and **T3**)
- A4:** *Could you identify a policy setting that is interesting to test in a full simulation run?* (**T1**, **T2**, and **T3**)

In the following, we report observations derived through a close reading and analysis of the notes gathered by the visualization researchers during the think-aloud case-study. We start with observations about general usage patterns, and afterwards we go through any remaining task-specific observations. Wherever relevant, we indicate when a visualization task (Sec. 3) is supported by the tool.

General Observations:

Exploring disease attributes – The mapping of different attributes to the individual nodes have been used regularly, and at different stages in the analysis. Infector state stood out as being the most frequently used, likely due to its relevance for understanding the impact of asymptomatic and presymptomatic transmission. For instance, during **A1**, **E2** commented "*Presymptomatics seem to be playing a disproportionate role in causing infections,*" and **E3** mentioned "*large trees have several asymptomatic nodes*" which were considered important to understand the behavior of the model (**T1**, **T2**). Location of infections was another factor that drew interest, and particularly the impact of school infections was initially commented as "*not seeming to jump out as massively important,*" but upon inspection of larger chains it was observed that school infections were featuring more, which prompted the experts to hypothesize that "*perhaps the school are important in the bigger chains?*" (**T2**, see also use case 2).

Degrees-of-freedom in interaction – The ability to compare multiple policy types, multiple attributes, and the impact of parameters, such as the *app-uptake*, was one of the key instruments that the experts used for comparing policies and evaluating the effectiveness of the models. However, with so many conditions to vary simultaneously, we observed that it became challenging to understand the impact of any factor, e.g., comparing two different policy types while also comparing different *app-uptake* values. As a result, the comparisons were mostly conducted with a single factor being varied with the other factors constant. For instance, during **A1**, setting

both sides to no-policy helped make the investigation of the infector state distributions much more straightforward.

Questioning simulation parameters – We observed that the visualizations were instrumental in exploring whether the parameters have the anticipated impact on the simulation outcomes. A good example was the *app-uptake* parameter that controls the percentage of the population using a contact-tracing application. While app-uptake was anticipated to have a high impact on the effectiveness of the policies during the modelling stage, it was discovered (during **A3** and **A4**, and as E1 commented) that "*(the impact of app-uptake is) definitely less dramatic than expected*" (**T1**). This was followed by setting the parameter to extreme values, i.e. 0% and 100% uptake. E2 commented that "*There is still an awful lot of green (prevented infections) on the tree, the impact seems to be coming from isolation instead of the app.*" After also exploring whether the low impact is due to the stringency of a policy by visually exploring a few stringent policies (**T3**), E1 concluded by stating "*it makes you wonder if app uptake is useful to include in the simulation.*"

Transmission trees – The variations within the topology and the patterns of transmission across various transmission trees have been one of the key foci of investigation. An immediate observation by E2, E3 and E4 during **A1** was that "*trees are quite short and not many trees go on for many generations*" (**T2**). During **A3**, such shorter chains have shown to respond differently to policies as E2 stated that "*not much happening with the short chains, as they happen quickly, not much time for policy to have impact.*"

Activity and Task Specific Observations:

While conducting A1, an initial strategy was to set a baseline and a significantly more stringent policy to compare. This was, however, found to be complicated when testing the validity of the model, so the analysis switched to a *same-policy* setup where a single variation was explored at a time. Node properties were most widely investigated during A1. In some cases, the exploration led to questions on the underlying data, for instance, when the nodes revealed very low infection levels at home, E2 asked whether this "*is this a glitch or just very few 'family'?*" (**T1**). After further inspection, there was indeed a glitch present where the 'family' location was being grouped into the 'other' location.

While conducting A2, a discussion on super spreading events, events where many transmissions happen in a short period of time, took place. E2 commented that "*there is no 'super spreader' built into the model, so not expected here, but there is a chance that such persons exist,*" and spotted a few likely cases with E1 making a comment that "*(they might be) a person that is highly connected*". E3 asked "*why are the trees so short? Expected them to be much longer*" and inquired the underlying disease parameters (**T1**, **T2**). A2 was also the only activity where the time element was actively explored, and it was not inspected extensively during the further activities related to the policies.

While conducting A3, one key analytical routine was to explore the relation between the strength of policies and their impact over different kinds of trees. For instance, E2 commented, while pointing at a large tree that, "*We are seeing on this part of the tree very little impact of the policy. [The infections] Manage to occasionally get a foothold. I guess you can see that there are branches where*

it [the policy] does not manage to curtail at all. Perhaps we compare with the feeble policy." (**T2**). And upon doing that, E2 further commented "*we can see the stringent policy have an impact*" (**T3**) where E1 hypothesized that "*maybe it is the time? We are not asking them to isolate very long or looking back.*" (**T2**).

While conducting A4, observations on the impact of application uptake parameter surfaced as a concern (as discussed above) and E1 suggested that it could potentially be left out of consideration. The investigation continued with a closer inspection of the x and y parameters, i.e., days of isolation and period of alerting, within this activity in order to propose a good trade-off between isolation time and outbreak control. E2 suggested "*14 days seem like a long time for isolation, let's choose 7 days*" (**T3**) and after seeing the results, stated that "*left and right not that different, so rather than 14 days, maybe 7 days is enough to have similar results.*" This followed on by a suggestion by E1 to compare 3 and 7 days, E2 commented that "*there is more difference between 7 days to 3 days*" (**T3**) prompting a decision on a good trade-off within the policy setting.

8.2. Semi-Structured Interview

Following the think-aloud case study exercise, a semi-structured interview of around 30 minutes took place. Discussion questions focused on eliciting reflections on four main themes: on the tool as a whole, on the individual designs, on the design process, and on wider and future applications of the designs.

New ways of seeing the simulation data – It took a while for the experts to orient themselves within the tool and E2 stated that "*(they) don't normally look at this kind of stochastic models in this way.*" However, as they went through the tasks, they became more comfortable formulating exploration routines. E2 acknowledged that by saying "*if we used the classic ways of looking at this data [referring to the aggregated disease curves], we wouldn't get the sense of what is happening in the infection clusters*" and "*would not have got that sense that sometime you completely nail clusters [with a policy] and that some other clusters are continuing.*" Similarly, E4 commented that they "*normally work with population models,*" and as also acknowledged by E3, that "*there is a lot of information in the trees*" and praised the potential usefulness of these in modelling processes.

Usability – Both groups of epidemiologists were able to use the tool without supervision quickly after working with it, requiring little prompting from the visualization experts. E3 and E4 however required additional explanation at the start due to unfamiliarity with the underlying model and network visualizations.

Representative Trees – The epidemiologists found the representative trees "*useful and intuitive*", and they allow for a large amount of trees that are not interesting individually to be summarized together. They did not feel like the representative trees overly summarize the data as they accurately capture the amount of nodes in the tree and as E2 remarked: "*it is possible to see the individual trees and how they are averaged*". This data provenance helped build confidence that the representative trees are good representatives.

Statistics and network measures – E3 and E4 expressed clear interest in the communication of statistics in relation to the simulations as a whole, and in relation to the individual trees. Although

visual summaries have been found to be useful, primarily due to the prevalence of quantitative measures in epidemiological modelling, statistical summaries and indicators could complement the visual patterns observed to build a well-rounded assessment of the policies. Similarly, there has been interest in making more use of network measures to organize the trees, e.g., ordering the representative trees by their depths to highlight long infection chains.

Emergency design process – The epidemiologists E1 and E2 found the emergency design process and being able to view the early visualizations overall useful to build a dialogue and gain insight. E2: "I think that [the emergency design process] was useful. Very useful when we saw the length of transmission chains: How that broke down under different policies. The idea culminated into the visualization. Even if we aren't ultimately using them, they created a dialogue about the kinds of things we look for in visualization and the kind of questions we ask."

Generalizability & future – While the current work is focused on visualizing a model for COVID data, E1, E2 and E3 all commented that there are "Definitely (other) settings where this can be useful to make suggestions to change policies," such as disease outbreaks in humans or livestock.

9. Discussion and Future Work

With our epidemiological experts, we made the decision to prune a single run of the simulation to allow for comparative analysis visualization of different policies (as opposed to having multiple runs for comparisons). This means that we are looking at a particular run from a stochastic model and not necessarily an average one. There are ways of mitigating this risk, for instance, creating ensembles and selecting the most representative/average/typical ones. Our approach only compares two policies and not multiple. Comparing multiple policies, possibly through visual parameter space analysis [SHB*14], would be interesting future work.

The simulation that is used to generate the data for our visualization builds on two different models: One model for mixing networks between different ages at different locations, and one model for baseline infectiousness at different locations. Both of these models have an inherent uncertainty within them, and it would be an interesting area of future work to understand and visualize the propagation of these uncertainties in the final visualizations.

While the tool is built for simulated data, it should be possible to adapt it to real infection maps from actual contact tracing data. One important difference is that in real data, often only potential avenues of infections are available instead of the exact infector-infectee relationships, but this could be resolved by using the work of Rozenstein et al. [RGPV16] to construct a likely infection map from the data.

In this paper, our approach selects representative trees according to specifics of the domain as discussed in Section 5.1. However, different design decisions can be made according to the area. Instead of calculating multiple representative trees for each group of trees of the same size, one can instead calculate a single representative tree to represent the entire cluster. This will remove the need for a user-specified distance measure as only a single tree per cluster

will be shown, and thus further reduce the number of trees shown on-screen. Furthermore, in the current approach a computationally expensive tree mapping algorithm is used to map a tree to its representative, which could be replaced with an inexact but faster mapping. Finally, the distance measure we are using is domain-specific and unlikely to work directly outside an epidemiological setting. However, any other distance measure can be used, and thus different domain-specific or even generic distance measures could be used for different application areas.

Lessons learnt – Reflecting on the discussions and evaluations, we distill some key lessons-learnt from this design study:

- Ability to vary data and conditions in comparative visual analysis is essential for evaluating multiple model outcomes but comparisons should not involve many varying factors at once
- Ability to concurrently compare the same data at multiple levels of aggregation is a key strength of visualization in this problem
- It is essential to use domain relevant metrics when constructing representatives. Using R_t values has been key in generating representatives that experts understand and trust
- Visualizing the internal dynamics of simulations enriches the analysis. Visualizing the infection maps through representative trees broadened the scope of analysis away from aggregate statistics and led to more nuanced discussions on comparing policies

10. Conclusion

In this paper, we present a visual analytics approach for visualizing different contact tracing policies on the results of individual-based model simulations for COVID-19. Our approach clusters transmission trees of similar epidemic structure together and uses representative trees to represent these clusters. The influence of different contact tracing policies and settings can be investigated on this infection map as a post-process to investigate effectiveness and facilitate model understanding. Through an evaluation with epidemiologists, we find that the approach helps understand the results, identify issues with the models, explore the impact of parameters, and compare and identify effective policies for further analysis.

With this study, we demonstrate that incorporating interactive visualizations within the various stages of the epidemiological model building process can lead to effective, better informed, and reflective model building processes, that contribute to both the quality of the models and the policies that are assessed through them. Visualization has more to offer to contribute to data and model intensive policy making scenarios, and strengthen the scientific communities' and society's resilience against the current and the future pandemics. To that end, we would like to stress, once again, the importance of close multidisciplinary collaborations in delivering effective and thoughtfully designed visual analytics approaches to support complex situations as in the pandemic response, and call for networks of interdisciplinary research teams to facilitate these.

Acknowledgments. We thank all our colleagues from the SCRC and RAMPVIS consortiums for input and feedback, as well as the RAMP initiative. This work was funded by the UKRI EPSRC grants EP/V033670/1 and EP/V054236/1.

References

- [AAP*05] ALLENDOERFER K., ALUKER S., PANJWANI G., PROCTOR J., STURTZ D., VUKOVIC M., CHEN C.: Adapting the cognitive walk-through method to assess the usability of a knowledge domain visualization. In *IEEE Symposium on Information Visualization* (2005), pp. 195–202. doi:10.1109/INFVIS.2005.1532147. 8
- [AFH*13] AKUTSU T., FUKAGAWA D., HALLDÓRSSON M. M., TAKASU A., TANAKA K.: Approximation and parameterized algorithms for common subtrees and edit distance between unordered trees. *Theoretical Computer Science* 470 (2013), 10–22. 6
- [AGJS*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* (2020), IEEE, pp. 86–90. 2
- [AMA21] ARLEO A., MIKSCH S., ARCHAMBAULT D.: A multilevel approach for event-based dynamic graph drawing. In *Eurographics/IEEE VGTC Conference on Visualization: Short Papers* (2021), Agus M., Garth C., Kerren A., (Eds.), pp. 103–107. doi:10.2312/evs.20211063. 3
- [APP10] ARCHAMBAULT D., PURCHASE H. C., PINAUD B.: Difference map readability for dynamic graphs. In *Proceedings of Graph Drawing* (2010), Springer-Verlag, p. 50–61. 2
- [Arc09] ARCHAMBAULT D.: Structural differences between two graphs through hierarchies. In *Proceedings of Graphics Interface* (2009), p. 87–94. 2
- [ASGK21] ANTWEILER D., SESSLER D., GINZEL S., KOHLHAMMER J.: Towards the Detection and Visual Analysis of COVID-19 Infection Clusters. In *EuroVis Workshop on Visual Analytics* (2021). doi:10.2312/eurova.20211097. 2
- [BDA*17] BACH B., DRAGICEVIC P., ARCHAMBAULT D., HURTER C., CARPENDALE S.: A descriptive framework for temporal data visualizations based on generalized space-time cubes. *Computer Graphics Forum* 36, 6 (2017), 36–61. doi:https://doi.org/10.1111/cgf.12804. 3
- [BHJ09] BASTIAN M., HEYMAN S., JACOMY M.: Gephi: An open source software for exploring and manipulating networks. *International AAAI Conference on Weblogs and Social Media* (2009). URL: <http://www.aaai.org/ocs/index.php/ICWSM/09/paper/view/154.3>
- [Bil05] BILLE P.: A survey on tree edit distance and related problems. *Theoretical Computer Science* 337, 1-3 (2005), 217–239. 6
- [BMC*20] BYRNE A. W., MCEVOY D., COLLINS A. B., HUNT K., CASEY M., BARBER A., BUTLER F., GRIFFIN J., LANE E. A., MCALOON C., O'BRIEN K., WALL P., WALSH K. A., MORE S. J.: Inferred duration of infectious period of SARS-CoV-2: Rapid scoping review and analysis of available evidence for asymptomatic and symptomatic COVID-19 cases. *BMJ Open* 10, 8 (2020). doi:10.1136/bmjopen-2020-039856. 2
- [BOH11] BOSTOCK M., OGIEVETSKY V., HEER J.: D³ data-driven documents. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (2011), 2301–2309. 6
- [BPW*21] BAUMGARTL T., PETZOLD M., WUNDERLICH M., HOHN M., ARCHAMBAULT D., LIESER M., DALPKE A., SCHEITHAUER S., MARSCHOLLEK M., EICHEL V. M., MUTTERS N. T., CONSORTIUM H., LANDESBERGER T. V.: In search of patient zero: Visual analytics of pathogen transmission pathways in hospitals. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2021), 711–721. doi:10.1109/TVCG.2020.3030437. 2
- [BWMM15] BRYAN C., WU X., MNISZEWSKI S., MA K.-L.: Integrating predictive analytics into a spatiotemporal epidemic simulation. In *IEEE Conference on Visual Analytics Science and Technology* (2015), pp. 17–24. 2
- [CAD*14] CARROLL L. N., AU A. P., DETWILER L. T., FU T.-C., PAINTER I. S., ABERNETHY N. F.: Visualization and analytics tools for infectious disease epidemiology: A systematic review. *Journal of Biomedical Informatics* 51 (2014), 287–298. 2
- [CARA*21] CHEN M., ABDUL-RAHMAN A., ARCHAMBAULT D., DYKES J., SLINGSBY A., RITSOS P., TORSNEY-WEIR T., TURKAY C., BACH B., BORGIO R., BRETT A., FANG H., JIANU R., KHAN S., LARAMEE S., NGUYEN P. H., REEVE R., ROBERTS J. C., VIDAL F., WANG Q., WOOD J., XU K.: RAMPVIS: Answering the challenges of building visualisation capabilities for large-scale emergency responses. <https://api.newton.ac.uk/website/v0/events/preprints/NI20011>, 2021. Accessed: 01-12-2021. 3
- [CBGM*21] CASEY-BRYARS M., GRIFFIN J., MCALOON C., BYRNE A., MADDEN J., MC EVOY D., COLLINS Á., HUNT K., BARBER A., BUTLER F., LANE E. A., O'BRIEN K., WALL P., WALSH K., MORE S. J.: Presymptomatic transmission of SARS-CoV-2 infection: a secondary analysis using published data. *BMJ Open* 11, 6 (2021). doi:10.1136/bmjopen-2020-041240. 2
- [DG14] DORIN A., GEARD N.: The practice of agent-based model visualization. *Artificial Life* 20, 2 (2014), 271–289. 4
- [DHA*20] DIXIT R. A., HURST S., ADAMS K. T., BOXLEY C., LYSEN-HENDERSHOT K., BENNETT S. S., BOOKER E., RATWANI R. M.: Rapid development of visualization dashboards to enhance situational awareness of COVID-19 telehealth initiatives at a multihospital healthcare system. *Journal of the American Medical Informatics Association* 27, 9 (2020), 1456–1461. 2
- [FWK*20] FERRETTI L., WYMANT C., KENDALL M., ZHAO L., NURTAY A., ABELER-DÖRNER L., PARKER M., BONSALE D., FRASER C.: Quantifying SARS-CoV-2 transmission suggests epidemic control with digital contact tracing. *Science* 368, 6491 (2020). doi:10.1126/science.abb6936. 2
- [GBA20] GÜRSAKAL N., BATMAZ B., AKTUNA G.: Drawing transmission graphs for COVID-19 in the perspective of network science. *Epidemiology & Infection* 148 (2020). 2
- [GCC*20] GRIFFIN J., CASEY M., COLLINS Á., HUNT K., MCEVOY D., BYRNE A., MCALOON C., BARBER A., LANE E. A., MORE S.: Rapid review of available evidence on the serial interval and generation time of COVID-19. *BMJ Open* 10, 11 (2020). doi:10.1136/bmjopen-2020-040263. 2
- [GHK10] GANSNER E. R., HU Y., KOBOUROV S. G.: Gmap: Drawing graphs as maps. In *Graph Drawing* (2010), pp. 405–407. 3
- [git] Contact-tracing-model. <https://github.com/ScottishCovidResponse/Contact-Tracing-Model.git>. Accessed: 22-11-2021. 2
- [GPLR21] GOMEZ J., PRIETO J., LEON E., RODRÍGUEZ A.: Infección—an agent-based model for transmission of infectious diseases: The COVID-19 case in bogotá, colombia. *PLOS ONE* 16, 2 (2021). 2
- [GXZ*18] GUO S., XU K., ZHAO R., GOTZ D., ZHA H., CAO N.: EventThread: Visual summarization and stage analysis of event sequence data. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2018), 56–65. 2
- [HACH12] HARRIGAN M., ARCHAMBAULT D., CUNNINGHAM P., HURLEY N.: Egonav: Exploring networks through egocentric specializations. In *Proceedings of the International Working Conference on Advanced Visual Interfaces* (2012), p. 563–570. doi:10.1145/2254556.2254661. 2
- [HB03] HARROWER M., BREWER C. A.: Colorbrewer.org: an online tool for selecting colour schemes for maps. *The Cartographic Journal* 40, 1 (2003), 27–37. 6
- [He20] HELLEWELL J., ET AL. S.: Feasibility of controlling COVID-19 outbreaks by isolation of cases and contacts. *The Lancet Global Health* 8, 4 (2020), e488 – e496. 2
- [HHO*16] HAKONE A., HARRISON L., OTTLEY A., WINTERS N., GUTHEIL C., HAN P. K., CHANG R.: PROACT: Iterative design of a patient-centered visualization for effective prostate cancer health risk communication. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2016), 601–610. 2

- [HPN*20] HINCH R., PROBERT W., NURTAY A., KENDALL M., WYMANT C., HALL M., LYTHGOE K., CRUZ A., ZHAO L., STEWART A., FERRETTI L., PARKER M. J., MÉROUEH A., MATHIAS B., STEVENSON S., MONTERO D., WARREN J., MATHER N., FINKELSTEIN A., BONSALE D., FRASER C.: Effective configurations of a digital contact tracing app: A report to NHSX. https://cdn.theconversation.com/static_files/files/1009/Report_-_Effective_App_Configurations.pdf?1587531217, 2020. 2
- [HS12] HOLME P., SARAMÄKI J.: Temporal networks. *Physics Reports* 519, 3 (2012), 97–125. 1
- [imp] Implementation of the tool. <https://github.com/maxie12/RepresentativeVis>. Accessed: 08-03-2022. 6
- [JPH*21] JAMES A., PLANK M. J., HENDY S., BINNY R., LUSTIG A., STEYN N., NESDALE A., VERRALL A.: Successful contact tracing systems for COVID-19 rely on effective quarantine and isolation. *PLOS ONE* 16, 6 (2021), 1–14. doi:10.1371/journal.pone.0252499. 2
- [Kar72] KARP R. M.: Reducibility among combinatorial problems. In *Complexity of Computer Computations*. Springer, 1972, pp. 85–103. 6
- [Ke20] KUCHARSKI A. J., ET AL. E.: Effectiveness of isolation, testing, contact tracing, and physical distancing on reducing transmission of sars-cov-2 in different settings: A mathematical modelling study. *The Lancet Infectious Diseases* 20, 10 (2020), 1151–1160. 2
- [KFS13] KOOP D., FREIRE J., SILVA C. T.: Visual summaries for graph collections. In *IEEE Pacific Visualization Symposium* (2013), pp. 57–64. 2
- [KOIY14] KONDO S., OTAKI K., IKEDA M., YAMAMOTO A.: Fast computation of the tree edit distance between unordered trees using ip solvers. In *International Conference on Discovery Science* (2014), pp. 156–167. 6
- [KRB*20] KRETZSCHMAR M. E., ROZHNVA G., BOOTSMA M. C., VAN BOVEN M., VAN DE WIJGERT J. H., BONTEN M. J.: Impact of delays on effectiveness of contact tracing strategies for COVID-19: a modelling study. *The Lancet Public Health* 5, 8 (2020), e452–e459. 2
- [KZA10] KOENIG P.-Y., ZAIDI F., ARCHAMBAULT D.: Interactive searching and visualization of patterns in attributed graphs. In *Proceedings of Graphics Interface* (2010), p. 113–120. 2
- [LBI*11] LAM H., BERTINI E., ISENBERG P., PLAISANT C., CARPENDALE S.: Empirical studies in information visualization: Seven scenarios. *IEEE Transactions on Visualization and Computer Graphics* 18, 9 (2011), 1520–1536. 8
- [Lew82] LEWIS C.: Using the 'thinking-aloud' method in cognitive interface design. *Research Report RC9265*, IBM TJ Watson Research Center (1982). 8
- [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* (2020), pp. 56–60. doi:10.1109/VIS47514.2020.00018. 2
- [LSDK18] LIU Y., SAFAVI T., DIGHE A., KOUTRA D.: Graph summarization methods and applications: A survey. *ACM Computing Surveys* 51, 3 (2018), 1–34. 2
- [MCH*20] MCALOON C., COLLINS Á., HUNT K., BARBER A., BYRNE A. W., BUTLER F., CASEY M., GRIFFIN J., LANE E., MCEVOY D., WALL P., GREEN M., O'GRADY L., MORE S. J.: Incubation period of COVID-19: a rapid systematic review and meta-analysis of observational research. *BMJ Open* 10, 8 (2020). doi:10.1136/bmjopen-2020-039652. 2
- [MDC*18] MOHR S., DEASON M., CHURAKOV M., DOHERTY T., KAO R. R.: Manipulation of contact network structure and the impact on foot-and-mouth disease transmission. *Preventive Veterinary Medicine* 157 (2018), 8–18. doi:https://doi.org/10.1016/j.prevetmed.2018.05.006. 2
- [MFM*18] MUELLNER U., FOURNIÉ G., MUELLNER P., AHLSTROM C., PFEIFFER D. U.: epidemix—an interactive multi-model application for teaching and visualizing infectious disease transmission. *Epidemics* 23 (2018), 49–54. 2
- [MGT*03] MUNZNER T., GUIMBRETIERE F., TASIRAN S., ZHANG L., ZHOU Y.: TreeJuxtaposer: Scalable tree comparison using focus+context with guaranteed visibility. In *ACM SIGGRAPH* (2003), pp. 453–462. doi:10.1145/1201775.882291. 2
- [MHW*21] MARVEL S. W., HOUSE J. S., WHEELER M., SONG K., ZHOU Y.-H., WRIGHT F. A., CHIU W. A., RUSYN I., MOTSINGER-REIF A., REIF D. M.: The COVID-19 pandemic vulnerability index (pvi) dashboard: Monitoring county-level vulnerability using visualization, statistical modeling, and machine learning. *Environmental Health Perspectives* 129, 1 (2021), 017701. 2
- [MLL*13] MONROE M., LAN R., LEE H., PLAISANT C., SHNEIDERMAN B.: Temporal event sequence simplification. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2227–2236. 2
- [MLR*11] MACIEJEWSKI R., LIVENGOOD P., RUDOLPH S., COLLINS T. F., EBERT D. S., BRIGANTIC R. T., CORLEY C. D., MULLER G. A., SANDERS S. W.: A pandemic influenza modeling and visualization tool. *Journal of Visual Languages & Computing* 22, 4 (2011), 268–278. 2
- [MPW*20] MÜLLER M., PETZOLD M., WUNDERLICH M., BAUMGARTL T., HÖHN M., EICHEL V., MUTTERS N., SCHEITHAUER S., MARSCHOLLEK M., VON LANDESBERGER T.: Visual analysis for hospital infection control using a rnn model. In *EuroVis Workshop on Visual Analytics* (2020), EuroGraphics, pp. 073–077. 2
- [NC09] NISHIURA H., CHOWELL G.: The effective reproduction number as a prelude to statistical estimation of time-dependent epidemic trends. In *Mathematical and Statistical Estimation Approaches in Epidemiology*. Springer, 2009, pp. 103–121. 4
- [PL20] PREIM B., LAWONN K.: A survey of visual analytics for public health. *Computer Graphics Forum* 39, 1 (2020), 543–580. 2
- [QCH*21] QUILTY B. J., CLIFFORD S., HELLEWELL J., RUSSELL T. W., KUCHARSKI A. J., ET AL. F.: Quarantine and testing strategies in contact tracing for SARS-CoV-2: a modelling study. *The Lancet Public Health* 6, 3 (2021), e175–e183. 2
- [RGPV16] ROZENSHTEN P., GIONIS A., PRAKASH B. A., VREEKEN J.: Reconstructing an epidemic over time. In *International Conference on Knowledge Discovery and Data Mining* (2016), pp. 1835–1844. 10
- [RPH*20] ROGERS J., PATTON A. H., HARMON L., LEX A., MEYER M.: Insights from experiments with rigor in an evobio design study. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2020), 1106–1116. 3
- [RPS*21] RIBEIRO P., PAREDES P., SILVA M. E., APARICIO D., SILVA F.: A survey on subgraph counting: Concepts, algorithms, and applications to network motifs and graphlets. *ACM Computing Surveys* 54, 2 (2021), 1–36. 2
- [RWA*13] RIND A., WANG T. D., AIGNER W., MIKSCH S., WONG-SUPHASAWAT K., PLAISANT C., SHNEIDERMAN B.: Interactive information visualization to explore and query electronic health records. *Foundations and Trends in HCI* 5, 3 (2013), 207–298. 2
- [SAK17] SIMONETTO P., ARCHAMBAULT D., KOBOUROV S.: Drawing dynamic graphs without timeslices. In *International Symposium on Graph Drawing and Network Visualization* (2017), Springer, pp. 394–409. 3
- [SAK20] SIMONETTO P., ARCHAMBAULT D., KOBOUROV S.: Event-based dynamic graph visualisation. *IEEE Transactions on Visualization and Computer Graphics* 26, 7 (2020), 2373–2386. doi:10.1109/TVCG.2018.2886901. 3
- [SBTK08] SCHRECK T., BERNARD J., TEKUSOVA T., KOHLHAMMER J.: Visual cluster analysis of trajectory data with interactive kohonen maps. In *IEEE Symposium on Visual Analytics Science and Technology* (2008), pp. 3–10. doi:10.1109/VAST.2008.4677350. 3

- [SHB*14] SEDLMAIR M., HEINZL C., BRUCKNER S., PIRINGER H., MÖLLER T.: Visual parameter space analysis: A conceptual framework. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 2161–2170. doi:10.1109/TVCG.2014.2346321. 10
- [SMM12] SEDLMAIR M., MEYER M., MUNZNER T.: Design study methodology: Reflections from the trenches and the stacks. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (2012), 2431–2440. doi:10.1109/TVCG.2012.213. 3
- [SMR*20] SYEDA U. H., MURALI P., ROE L., BERKEY B., BORKIN M. A.: Design study "lite" methodology: Expediting design studies and enabling the synergy of visualization pedagogy and social good. In *Proceedings of Conference on Human Factors in Computing Systems* (2020). URL: <https://doi.org/10.1145/3313831.3376829>. 3
- [STM*22] SONDAG M., TURKAY C., MOHR S., MATTHEWS L., XU K., ARCHAMBAULT D.: What is R? A graph drawer's perspective. In *International Symposium on Graph Drawing and Network Visualization* (2022), Springer. 4
- [vLGS09] VON LANDESBERGER T., GÖRNER M., SCHRECK T.: Visual analysis of graphs with multiple connected components. In *IEEE Symposium on Visual Analytics Science and Technology* (2009), pp. 155–162. doi:10.1109/VAST.2009.5333893. 3
- [WHA07] WILLETT W., HEER J., AGRAWALA M.: Scented widgets: Improving navigation cues with embedded visualizations. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1129–1136. doi:10.1109/TVCG.2007.70589. 6
- [WXL*21] WANG Y., XIONG H., LIU S., JUNG A., STONE T., CHUKOSKIE L.: Simulation agent-based model to demonstrate the transmission of COVID-19 and effectiveness of different public health strategies. *Frontiers in Computer Science* (2021). 2
- [YDH*17] YANEZ A., DUGGAN J., HAYES C., JILANI M., CONNOLLY M.: PandemCap: Decision support tool for epidemic management. In *IEEE Workshop on Visual Analytics in Healthcare* (2017), IEEE, pp. 24–30. 2