Workshop on Visualisation in Environmental Sciences (EnvirVis) (2024)
S. Dutta, K. Feige, K. Rink, and B. Nsonga (Editors)

WeaVA: Weather Event Characterization Based on Citizen Reports

Dominique Haessig[1], Haiyan Yang[2], Renato Pajarola[2], Alexandra Diehl[2]

[1]MeteoSwiss, Switzerland, [2]University of Zürich, Switzerland

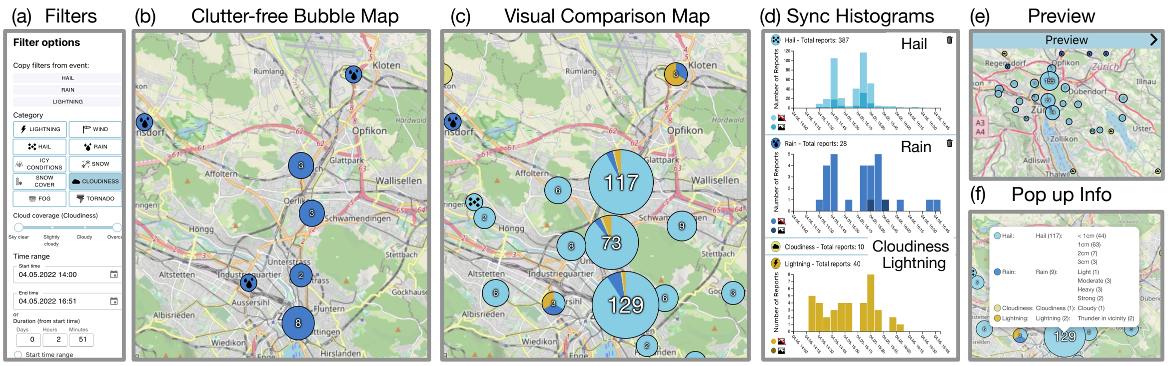

Figure 1: WeaVA *is a visual analytics solution for the interactive exploration and analysis of weather observations as reported by citizens. Our visual design combines flexible filters with optimized clutter-free maps, and synchronized histograms for an effective analysis of citizens' weather reports, characterization of high-impact and unexpected weather events, and comparative analysis of multiple events.*

Abstract

Observations of weather captured by citizens represent a novel and unique data source that can complement other authoritative sources, such as remote sensing, and help detect and characterize high-impact weather events. This work proposes a visual tool that characterizes weather events by visually analyzing online citizens' reports gathered by MeteoSwiss, the Swiss Federal Office of Meteorology and Climatology. Our solution supports the visual exploration of selected features like weather event categories and intensities through time and space. It presents a novel clutter-free bubble map visualization that facilitates an easy exploration and quantification of weather reports. It allows the analysis at different zoom levels, supporting multiple interactive exploration features such as synchronous or asynchronous event histogram comparisons, clutter-free pie-chart map visualizations, and animations. We illustrate our approach with a series of use cases and findings. We performed a user study with domain experts from the national weather services in Switzerland, Austria, and Argentina to evaluate our tool's expressiveness, effectiveness, and easiness of use. We also list the benefits of our design, future work, and limitations.

CCS Concepts

• ***Human-centered computing*** → ***Visualization design and evaluation methods; Visual analytics;*** *Geographic visualization;*

1. Introduction

Weather analysis and prediction is a large and longtime field of interest for humans. Many methods and sources are used to gain information about the weather, such as monitoring stations, radio soundings, or satellites. To improve the analysis and predictions, new additional ways are searched. One of these newer ways is the gathering of data with the help of citizens.

Online human digital traces (text, image, video) about weather events are a unique and novel resource to support the understanding and characterization of high-impact weather events. There are several efforts from European Weather Services to collect weather event reports generated actively by citizens, for example the mPing app from the Deutscher Wetterdienst in Germany [Wet], or the European Severe Weather Database, short ESWD [Lab18], among others.

DOI: 10.2312/envirvis.20241136

In November 2021, MeteoSwiss launched a new feature, the Meteo reports, in their weather app, the MeteoSwiss App, in which citizens can report the weather conditions in their surrounding areas, optionally including images. Such reports are divided into categories like cloudiness, rain, hail, and wind, specified and provided by MeteoSwiss. For each of these categories, there are degree levels for the respective intensities, like strength or dimension. Posted reports and optional images are checked for correctness and appropriateness before being included in the app. As a first step to gaining more knowledge on how citizens report weather events, we investigate the following main research questions (RQs):

R1: What weather event features are reported by citizens, and how are they linked through the reports?
R2: How do people report high-impact weather events over time and space?
R3: What patterns or unexpected activity can be extracted from the visual inspection of weather reports?

To explore these RQs, we propose a visual analytics tool that facilitates the visual comparison of weather reports and their underlying weather events (see Fig. 1). Our tool presents the following main contributions:

- A clutter-free bubble map visualization of the citizens' reports with comparative features. Our approach combines a fast bubble map generation with a manageable amount of visual marks on maps representing the locations. The shape of the bubbles is used for additional information representation.
- An synchronous/asynchronous (sync/a-sync) histogram view that facilitates the comparison of multiple categories and similar weather events over time. We refer to *sync* histograms to the histogram views that are either synchronized on the starting date and time or by duration. This view, combined with a time player and linked to the map, allows the user to follow the reports' appearances and changes in time and space.
- Multiple features and flexible parametrization that allow the user to inspect potential patterns and unexpected events on the citizens' reports.

We present representative use cases showing how the tool can be utilized and what information can be gained. We evaluated our tool with seven domain experts to identify its expressiveness, ease of use, and effectiveness, as well as possible improvements and limitations. We summarized the lessons learned from using our solution and the feedback from the user study.

The novelty of our work consists of (1) a balanced design that combines spatial and temporal views with a flexible set of parameters and filters to support the user in the characterization of weather events through citizens' report properties, (2) the identification of spikes of attention in particular severe event cases, (3) the analysis of the spatio-temporal citizens' contributions, and (4) the multivariate analysis and verification of uncommon reports.

2. Related Work

2.1. Weather Visualization Tools

Weather visualization is a longstanding active research topic. Rautenhaus et al. [RBS*18] covered the related work in visualization tools used for meteorological data analysis. More specific tools for nowcasting and warning of severe weather hazards include AWIPS [UCAa] (the Advanced Weather Interactive Processing System), a meteorological decoding, display, and analysis package initially developed by the United States National Weather Service. WarnGen [UCAb] is a tool based on the AWIPS CAVE platform for creating and issuing weather warnings. WarnGen allows for the visualization of weather radar data. However, to the best of our knowledge, it does not provide information about crowdsourcing information or distribution maps of citizens' observation data sets.

2.2. Visual Analysis of Crowdsourced Data

We focus on a newly available dataset of crowdsourced weather reports collected by MeteoSwiss with unique features such as weather categories, intensities, and pictures taken by citizens. Similar efforts to collect citizen data are being made by the European Severe Weather Database `(eswd.eu)`, the "Unwetterzentrale" `(uwz.com)`, the "Zentralanstalt für Meteorologie und Geodynamik (ZAMG)" in Austria, the "Deutscher Wetterdienst "`(dwd.de)` in Germany, and the "Weather Observations Website" (WOW) of the Met Office in the United Kingdom"`(https://wow.metoffice.gov.uk/)`. These datasets have shown potential to be used as complementary data sources to study weather phenomena [CGH*16, BHM*19].

A close antecedent to our work is CitymisVis [HMCnM17], which analyzes requests and complaints about municipal services. In *CitymisVis*, the data is divided into categories such as street condition, running water, parks and squares, and corresponding subcategories. The visualization utilizes the geo-information of the report on a map and then uses a heat map and hierarchical clustering to highlight specific areas. Furthermore, detailed information on the reports in certain highlighted areas is shown in an additional pie chart where individual reports can be inspected. Besides the map, they provide statistics over the data set, like our additional information in the form of histograms. For them, the proportion of the categories and sub-categories is more relevant than the timing of the reports, so they decided to use a sunburst chart, where the inner ring is about the categories and the outer ring is about the subcategories. Also, De Melo Borges et al. proposed *EstaVis* to cluster and analyze reports of problems in urban infrastructure to facilitate the data overviews [DMBBP*16]. The main advantage of this work is that it can deal with streaming processing of incoming reports by a fast clustering computation mechanism.

CitizenSensing [NOV*20, NVO*21] is one of the most relevant to our work as it also focuses on weather conditions. It visualizes the report data in various forms: a map view for displaying the spatial locations of the report, a Sankey diagram showing the linkages between weather event type, impact type, and personal level of comfort, a word cloud view displaying frequent terms submitted in the comments of the reported weather events, and a temporal scatter plot, where each point corresponds to a report and is color-coded by level of comfort or climate impact. Besides, the application further provides temperature, air pressure, and humidity information that is not based on the weather reports but on a sensor network. Although providing multiple features and functionality, CitizenSensing has a limitation in that it does not scale well with a large number of reports, as the maps would be overfilled with location information, and the scatter plot would also be cluttered with too many points, whereas our work specifically addresses this issue.

2.3. Clutter-Free Spatial Data Visualization

An important point we needed to address in our work is how to visualize overplotted visual elements on the maps. We hereby summarize some of the most commonly used approaches for encoding spatial point data over 2D maps.

Dot plots or scatter plots, with various shapes, can represent data items on maps [ZKT21, WCD*19], but lack the ability for visual abstraction of potential clusters or dense data regions if showing the clusters or hot spots of data points is essential. Pixel maps [Kei00] are helpful, but their neighborhood pixel expansion criteria might mislead the visual interpretations of the phenomena. On the other hand, density plots, including heatmaps, can show different clusters but lose the information of individual data points [WXS12, ZLG*21]. Other antecedents include different line plots on maps, but they mainly deal with certain spatial information such as trajectories [KJW*18, ZCL*21, ARH*15].

Extending basic plots into well-designed glyphs also has a broad application in spatial data visualization [BKC*13, ML19]. Bergenstråhle et al. [BBL20] facilitate cluster evaluation with their R-based package SpatialCPie. It uses pie charts to indicate the similarity between spatial regions and clusters and further plots an array of such pie charts to show the similarity scores between each spot and the centroid of the analyzed cluster. Kumpf et al. [KTB*18] and Weng et al. [WXS12] both use a pie chart-like glyph design to represent cluster information in their applications, but the size of those pie chart glyphs is set to be a constant. The work of Rau et al. [RHWS22] presented different gridification methods to make the data clutter-free. McNabb and Laramee [ML19] used level-of-detail scale-aware glyph maps to reduce clutter. These two last antecedents are the closest to our clutter-free bubble method.

3. Methodology

In this section, we describe the participatory design process, the data model, and the visual features of our solution.

3.1. Data Sets

The main data set used for our solution is citizens' weather reports contributed by users of the visual tool named *MeteoSwiss App*. This data set consists of 501,053 reports. The time range of the data is in total from 10th of June 2021 to 6th of December 2023. We also use auxiliary data from the Federal Office of Topography SwissTopo to identify cities and towns, including postal codes and perimeter. The data is used for information about the cantons and their locations.

Reasons for reports not passing the quality check are failing validation with the weather data, being outside of the time range or supported region of the feature, or problems with the user like too many images per hour or lost trust. The reasons for removing images are that they contain humans, text, hate symbols, alcohol or drug products, or that there were problems with the upload. The filter is automated and removes images that are unsure so that no problems should be in the approved images.

3.2. Design Process

The visual design was done using an iterative and participatory design process. Together with experts from the MeteoSwiss, we have identified the main objectives and tasks of our tool:

T1: To visualize the spatio-temporal distribution of weather reports over time and space.
T2: To characterize weather events using multiple weather report categories.
T3: To visually compare two or more weather events identified by the users using different temporal ranges, regions, categories, and intensities.

Starting from an initial design prototype, we iterated over it with our collaborators in several rounds. In particular, their feedback helped us to add the required flexibility to the tool by expanding the parameters and feature filters, adding multiple selections of categories, and supporting the visual comparison of multiple events. In the last step of our design process, we included a broader audience of experts and computer scientists. The main feedback we received from this iteration was the need to compare several categories and intensities distributed in the map. To respond to this feedback, we designed the *clutter-free bubble map* that allows for a non-occluded view of the report distribution over time and space, distinguishing different categories per location employing pie charts.

3.3. WeaVA Visual Design

Our solution provides three main interactive features: (1) a stratified clutter-free bubble map visualization that facilitates a clutter-free exploration of a large number of citizen reports, and that allows the user at the same time to have an overview of the density as well as the distribution of reports over the map *(task 1 and 2)*, (2) a comparative visualization of weather events and weather event features *(task 3)*, and (3) a sync/a-sync event histogram view that facilitates the temporal comparison of events over time, synchronized by date, by duration, or without synchronization using independent time ranges *(task 1 and 2)*.

3.3.1. Clutter-free Bubble Map

There are different ways to present location information on a map, such as choropleth, markers, and density maps, among others. We have chosen to represent the report data with markers, as markers show the location they belong to with more precision and allow for the addition of several visual encodings into the marker.

We have also evaluated different basic clustering algorithms for aggregated bubble maps, but none produced satisfying results. K-Means clustering, for example, will generate different maps in each generation when the starting points are not set to be the same. DBSCAN produces stable results, but for large numbers of reports, typically, very dense maps are produced, leading to large areas that are difficult to connect to a single marker. Furthermore, DBSCAN areas can be non-convex so that there are points between a marker and a report that are not represented by the marker.

To solve the above mentioned issues, we apply a stratification method that divides the map into a uniform grid of a given size and computes the reports' distribution. For every grid cell, the reports belonging to it are aggregated into one marker, as illustrated in Fig. 2. The basic grid distance uses a predefined value and adjusts according to the zoom level on the map. Therefore, for each step of zooming-in, the distance covered per pixel halves and, accordingly, the distance of the zoom halves as well. Our method first calculates the actual grid size by dividing the total distance in each direction through the predefined maximal grid distance. The rounded-up grid distance then divides the total distance to get the size of one cell per direction. The grid will be initialized by assigning all reports to their respective cells, as shown in Fig. 2(b). Next, the convex hulls and the centroids of the contained reports are calculated for each cell, see Figs. 2(c) and 2(d). The so-generated centroids will be the coordinates of the bubble markers; see Fig. 2(e). Finally, different markers are created from spatially separated points, only slightly overlapping at the borders. We eventually use varying-sized circles as can be seen in Fig. 2(f), i.e., *bubbles*.

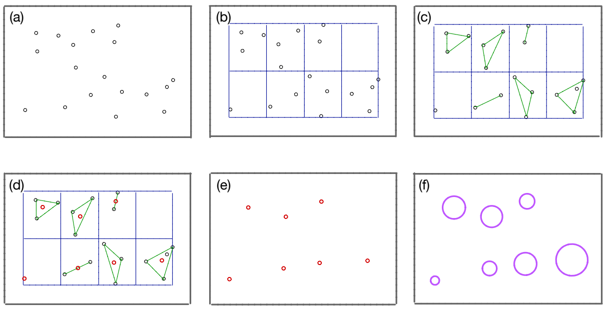

Figure 2: *Clutter-free Bubble Map algorithm steps: (a) initial report distribution, (b) grid initialization, (c) convex hull definition, (d) centroids extraction, (e) centroid locations, (f) clutter-free bubbles.*

A pie chart marker is displayed if there are multiple reports in one marker combined, whether it is because there are multiple reports at the same coordinate or clustered together. The size of each slice corresponds to the percentage reports of a respective event in relation to the total number of reports. Each pie chart slice's color corresponds to the respective event's color. We chose pie charts as a marker to show multiple categories in one location because they allow us to maintain a coherent design for the user and evoke the whole-and-part visual metaphor. Moreover, the center of a circle can more easily be connected to a spatial location than, for example, a bar. Combined with the different circle sizes, the user can have a rough qualitative estimate of the relative number of reports for the given event.

Another advantage of our method is that the average complexity of the algorithm is linear. It only loops once through the list of coordinates, which is the complexity for the best clustering algorithms.

The bubbles sometimes slightly overlap. This is due to the bubble center not being in the center of the grid cell but the points' convex hull's centroid. This is a trade-off with the information gain of the centroid. The centroid shows where the reports are more representative. Fig. 3 shows an example of a Clutter-free bubble map and its counterpart, a bubble map with clutter for many reports.

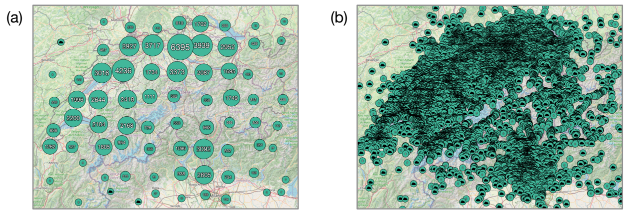

Figure 3: *Clutter-free Bubble Maps. (a) shows the results of the clutter-free bubble maps, and (b) shows how reports are depicted by using only location markers.*

The marker size is used as another visual variable to show the number of reports per marker in the cluster-free bubble map. We use a linear normalization, where the *rate* per cell is calculated as:

$$\text{rate}_i = \frac{\text{all reports in cell } i}{\text{max number reports in any single cell}} \quad (1)$$

Using this per-cell normalization rate, the size of a marker in its i-th cell is then calculated as:

$$\text{size}_i = s_{\text{min}} + (s_{\text{max}} - s_{\text{min}}) * \sqrt{\text{rate}_i} \quad (2)$$

where s_{min} and s_{max} are the minimal and maximal sizes of markers, and i is the cell in which the marker is located. The maximal size of a marker is the grid cell size. The root is taken for the size since a linear growing area is desired instead of a linear growing diameter.

3.4. Visual Comparison Map

The visual comparison map, as shown in Fig. 1(c), allows the user to compare different categories associated with an event. For example, reports of a thunderstorm might be related to wind, hail, and rain. Moreover, comparisons of similar events that happen at different time ranges or geographic locations can also be made. The second type of comparison is especially useful to characterize spatio-temporal changes or similarities of high-impact weather events.

3.5. Sync/A-sync Time Histogram View

The sync/a-sync histogram view plays an important role in the visual comparison of similar weather events, as reported by citizens. This temporal analysis and comparison view supports (1) the analysis of a particular weather event evolution in time and (2) the comparison of similar weather events that may occur at different times or locations. The synchronization can be done by the start date or the duration of the event or can be omitted, i.e., asynchronous, as further described below. The combined analysis of this view and the visual comparison map allows the user to have a global overview of similarities among multiple events of interest.

The histograms provide information about the number of reports with or without images over a selected time period of the event. The

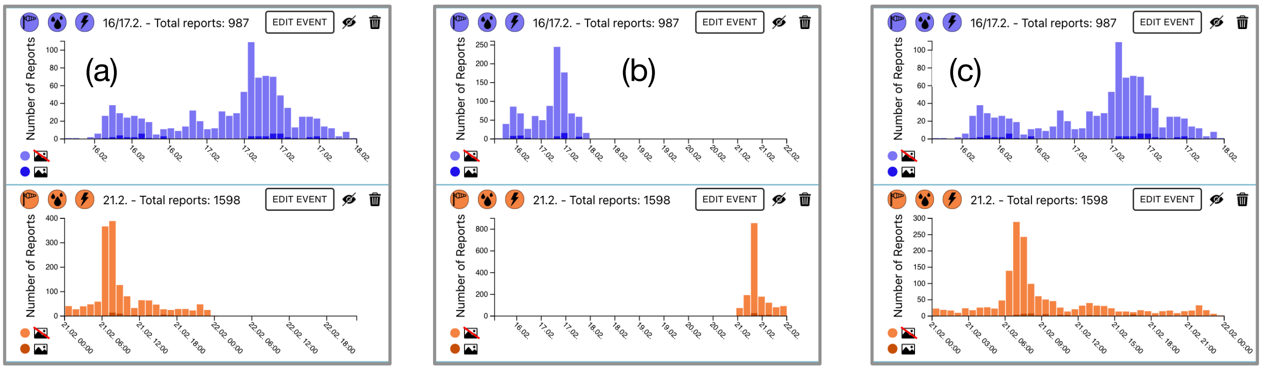

Sync Duration | Sync Duration and Start Time | No Sync

Figure 4: *The sync/a-sync histogram views: (a) shows two histograms corresponding to two different storm events on the 16th/ 17th of February and 21st of February synchronized by duration. (b) shows the same selected events synchronized by start and end date. (c) shows both events in the sync histogram views without synchronization.*

start and end times of the event are defined by the user. The lower and darker part of the histogram represents the reports with images, and the upper part represents the reports without images. The user can adjust the number of bins displayed to analyze the temporal distribution of the reports at different levels of detail. By default, the number of bins is set to 20, following Sahann et al. [SMS21], who found that 20 bins is the optimal number because the error rate of human observers becomes stable and does not improve with additional bins. However, the user is free to change this number as desired.

In the comparison view, the user can change the synchronization of the display of the different histograms. The options are:

Sync the duration: This option synchronizes the temporal axis of the histogram to the duration of the longest time range. The temporal axis of each event begins at its own start time but lasts the same as the longest time range. This option is best suited to compare events of the same category on different dates (see Fig. 4(a)).

Sync the start: This option synchronizes the start of the histogram temporal axis to the earliest start time. This option is best suited to compare events of different categories or intensities around the same time (see Fig. 4(b)). The events' duration is set to last the same as the longest time range of both events for comparability reasons.

Sync nothing: In this option, the histograms are not synchronized, neither start and end times nor duration of the events. The used boundaries are the ones selected in the selection view (see Fig. 4(c)). This option is useful when the users want a quick, broad search to spot interesting events.

The users can drill down into a temporal analysis by brushing and selecting a sub-range in the histograms. The user can select a subset of reports by clicking and dragging over multiple bins. The reports of the selected bins will simultaneously be highlighted on the map.

4. Use Cases

In this section, we describe two use cases to illustrate the main use, features and benefits of using our tool.

4.1. A Major Hail Event near Zurich

In this case, we analyze an event that took place in the afternoon of 4th of May 2022 in the city of Zurich. The event surprised the population and many citizens reported it, as portrayed in the news [Kle22]. The user selects *hail* as the leading category, the date of the event, and the preferred area and surroundings: i.e. the canton of Zurich. By visualizing the information in the sync/a-sync histograms, she narrows down the event to the approximated time of occurrence as shown in Fig. 5(a). Subsequently, she adds the cloudiness, rain, and lightning categories as shown in Fig. 5(b). Further, she slightly varies the time range to add further observations that were reported right before or after the peak time, as observed in the histograms (see Fig. 5(c)). The results of her analysis reveal several findings:

1. The number of reports per category varies a lot. The majority of the reports are about hail (387), followed by lightning (40), rain (28) and finally cloudiness (11). It shows that citizens report the most impactful category in case of a special weather event. These findings are aligned with R1 (Section 1) and T1 and T2 (Section 3).
2. An analysis of reported intensities shows high variability in the reports. For hail, the maximal number of reports is for a size of 1 cm, but sizes up to 3 cm are reported as well in the same reports group. These findings are aligned with R3 and T2.
3. The number of reports per area varies along the path of the weather front. Most reports came from areas around university buildings. These findings are aligned with R2 and T1 and T2.
4. The spatio-temporal analysis shows that a clear progress of the weather front can be seen, as citizens in the south report earlier than citizens in the north. These findings are aligned with R3, T1, and T2.

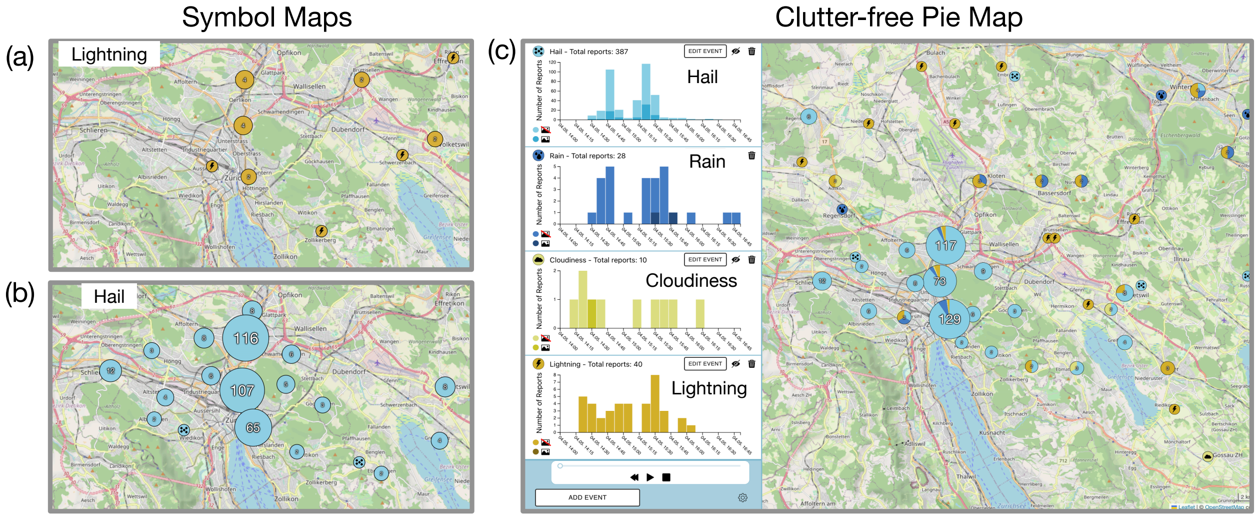

Figure 5: *Major hail event on 4th of May 2022, in the city of Zurich. (a) shows citizens' reports on the map corresponding to lighting, (b) shows citizens' reports on the map corresponding to hail, and (c) the visual comparison map showing how hail, rain, cloudiness, and lightning contribute to the specific high-impact weather event.*

These findings show that our tool facilitates a quick spatio-temporal overview of the number of reports for the same high-impact event. This information can be useful to study how citizens report weather events, what categories or kind of events are important for them to report, and how reports progress in time and space.

4.2. Comparison of two Similar Stormy Events

High-impact weather events repeat with certain periodicity over time and space, some of them are particularly local and vary slightly in their spatio-temporal extension. These characteristics make the analysis of similar or *analog* events of high interest not only for meteorologists but also for the general public.

To illustrate how our tool facilitates such an analysis of analog events, the user examines reports of heavy storm events and compares their similarities and differences. We chose an event of heavy storms that occurred in February 2022. The selection was made based on information from the climate bulletin of MeteoSwiss [Met22].

The user starts her exploration by selecting the category "wind", which is the main category associated with heavy storms. Then, she narrows down her analysis to high-intensity winds, so that no reports of breezes are included. With the help of the histogram, the user observes that, indeed, February 2022 is reported as the stormiest month in the available time range. The histograms show two peaks in February corresponding to the storms that took place on 6th and 21st February, see Figs. 6(a) and 6(b) respectively. The analysis of the two storm events reveals that:

1. The storm on 6th February shows more citizens' reports than the storm on 21st. However, the second storm has more reports in the Alps and in the South as shown in Fig. 6(c). (R1 and T3)
2. The analysis of the intensities shows that citizens reported it quite differently. Most citizens reported stormy winds, some even storms, and a few citizens reported the intensity of a hurricane, see Fig. 6(d). (R2, R3, and T2)
3. The spatio-temporal comparison done by playing through the reports of both events shows no difference in the spatial distribution. The number of reports grows similarly in the whole area. (R1, R2, and T1)
4. The storm in the morning shows that citizens reported the weather throughout the whole night. Even so, the number of reports strongly increased from 6:00 AM, a hypothesis could be that many citizens were awake. After 8:00 AM, the number decreased again rapidly, see Figs. 6(e) and 6(f). (R3 and T1)
5. The number of reports of the evening storm has a more continuous progression. The reports present two waves, with a local minimum around 19:00, which can be due to a break of the wind or other citizens' activities, for example, citizens eating at that time, see Figs. 6(e) and 6(f). (R3 and T1)
6. The spatio-temporal analysis of both events through the player shows that both were reported over the whole time in the lowlands of Switzerland. Unlike the hailstorm of the first use case, this event has no visible front. (R1, R2, R3, and T1)
7. The storm happened at different times of the day. The storm on 6th of February was in the early morning, the one on the 21st in the evening, as shown in Fig. 6(e). (R1, R3, and T3)

These findings raise new research questions about how some daily routines might be reflected in the data and how such phenomena might attract users' attention and interest in observing and reporting more similar events in the future.

5. Domain Expert Evaluation

We evaluated our visual design using a round of semi-structured interviews. Our interviewees consist of a diverse group of meteo-

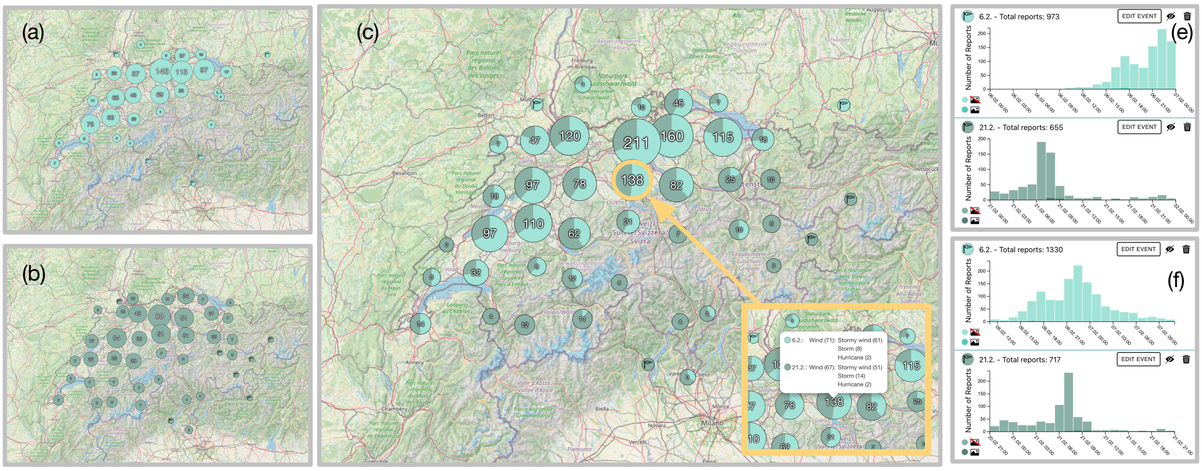

Figure 6: *Heavy storms registered during February 2022. (a) shows the distribution of wind reports on 6th of February, (b) the distribution on 21st of February. For both, the clutter-free bubble map is used. (c) shows the visual comparison of these two heavy storms. The pie chart markers show on one hand how prominent citizens' reports were on the upper north of the country, and second that in the same area people registered or noticed more the first event. (d) shows the detailed information of a bubble. Citizen reported different intensities, even on a hurricane level. (e) shows the events over the whole day of the respective event, emphasizing the day time of the storms, and (f) shows the same event with slightly shifted time ranges to align the peaks of the storm reports. There are only few reports containing images, which can be due to the nature of the category* wind*, which is difficult to capture in a static image.*

rologists, environmental scientists, and computer scientists working at the national weather services of Switzerland (MeteoSwiss), Austria (ZAMG), Europe, and Argentina in South America. We interviewed a total of seven participants. Five interviewees have more than ten years of experience in weather forecasting, and the other two have data science jobs directly linked to weather forecasting.

The interview contained a pre-experiment questionnaire, a pair analytics session, and a post-experiment questionnaire, and lasted approximately 45 minutes approximately. We provide the evaluation form as supplementary material.

Participants found the novel *clutter-free bubble map* the most effective among all the visual features. In general, all the features are evaluated as very effective, besides the preview lens which was rated moderately to very effective. They found the *sync histograms*, the *clutter-free bubble map*, and the *preview lens* moderately to very easy to understand, compared to other features. These results make sense because these visual features are enhanced for visual comparison tasks which makes them more complex. We observed a similar pattern in their responses regarding the feature expressiveness, but the participants' feedback was also very positive. They found all visual features very expressive to extremely expressive.

They gave positive feedback and mentioned a wide range of tasks that can already be addressed by our tool such as:

- identifying high-impact weather events,
- characterizing weather events' spatial distributions,
- performing evaluation and comparison of past weather events,
- complement the verification process of severe thunderstorm warning,
- analyzing citizens' activity when warnings are issued or when there is a missing warning of high-impact weather events.

They also mentioned possible future extensions for the tool such as spatial statistical reports, overlay of remote sensing observations and warning maps, the inclusion of reported images and filtering based on image feature extraction, export of customized reports for further analysis outside the tool, and real-time performance. The current code of our tool is publicly available at: https://github.com/dhaess/WeaVA.

6. Conclusions

In this paper, we introduced *WeaVA*, a visual analytics tool for characterizing weather events as reported by citizens. Our visual design and solution successfully answered our initial research questions in Sec. 1 and tasks in Sec. 3, as shown through the different use cases.

During the expert user evaluation, we identified further potential features and future work, such as the query by automated features extracted from images attached to the reports, the inclusion of remotely sensed data coming from radars and satellite images, and the possibility of importing data sets from different countries, in particular, the so-called European *DACH* area composed by Germany, Austria, and Switzerland.

A limitation is that our solution currently needs to support real-time analysis. To include them, we foresee using progressive visual analytics and high-performance computing methods.

7. Acknowledgements

The Swiss National Science Foundation supports this research through the SINERGIA grant for the project: "UnWeather Vizard: Uncertainty Visualization and Analysis of High-Resolution Numeric Weather Forecasts" (grant number: 208541).

References

[ARH*15] AMINI F., RUFIANGE S., HOSSAIN Z., VENTURA Q., IRANI P., MCGUFFIN M. J.: The impact of interactivity on comprehending 2d and 3d visualizations of movement data. *IEEE Transactions on Visualization and Computer Graphics 21*, 1 (2015), 122–135. doi:10.1109/TVCG.2014.2329308. 3

[BBL20] BERGENSTRÅHLE J., BERGENSTRÅHLE L., LUNDEBERG J.: SpatialCPie: an R/Bioconductor package for spatial transcriptomics cluster evaluation. *BMC Bioinformatics 21*, 1 (2020), 161. URL: https://doi.org/10.1186/s12859-020-3489-7, doi:10.1186/s12859-020-3489-7. 3

[BHM*19] BARRAS H., HERING A., MARTYNOV A., NOTI P.-A., GERMANN U., MARTIUS O.: Experiences with> 50,000 crowdsourced hail reports in switzerland. *Bulletin of the American Meteorological Society 100*, 8 (2019), 1429–1440. 2

[BKC*13] BORGO R., KEHRER J., CHUNG D. H., MAGUIRE E., LARAMEE R. S., HAUSER H., WARD M., CHEN M.: Glyph-based visualization: Foundations, design guidelines, techniques and applications. In *Eurographics (state of the art reports)* (2013), pp. 39–63. 3

[CGH*16] CHEN S., GOURLEY J. J., HONG Y., CAO Q., CARR N., KIRSTETTER P.-E., ZHANG J., FLAMIG Z.: Using citizen science reports to evaluate estimates of surface precipitation type. *Bulletin of the American Meteorological Society 97*, 2 (2016), 187 – 193. URL: https://journals.ametsoc.org/view/journals/bams/97/2/bams-d-13-00247.1.xml, doi:10.1175/BAMS-D-13-00247.1. 2

[DMBBP*16] DE MELO BORGES J., BUDDE M., PETERS O., RIEDEL T., SCHANKIN A., BEIGL M.: EstaVis: A real-world interactive platform for crowdsourced visual urban analytics. In *Proceedings on IoT in Urban Space* (2016), p. 65–70. URL: https://doi.org/10.1145/2962735.2962749, doi:10.1145/2962735.2962749. 2

[HMCnM17] HUBERT R. B., MAGUITMAN A. G., CHESÑEVAR C. I., MALAMUD M. A.: CitymisVis: A tool for the visual analysis and exploration of citizen requests and complaints. In *Proceedings on Theory and Practice of Electronic Governance* (2017), p. 22–25. URL: https://doi.org/10.1145/3047273.3047320, doi:10.1145/3047273.3047320. 2

[Kei00] KEIM D. A.: Designing pixel-oriented visualization techniques: Theory and applications. *IEEE Transactions on visualization and computer graphics 6*, 1 (2000), 59–78. 3

[KJW*18] KIM S., JEONG S., WOO I., JANG Y., MACIEJEWSKI R., EBERT D. S.: Data flow analysis and visualization for spatiotemporal statistical data without trajectory information. *IEEE Transactions on Visualization and Computer Graphics 24*, 3 (2018), 1287–1300. doi:10.1109/TVCG.2017.2666146. 3

[Kle22] KLE: Hagelgewitter zieht über Zürich – «so viel Eis habe ich noch nie gesehen». *20 Minuten* (2022). Accessed: December 2nd, 2022. URL: https://www.20min.ch. 5

[KTB*18] KUMPF A., TOST B., BAUMGART M., RIEMER M., WESTERMANN R., RAUTENHAUS M.: Visualizing confidence in cluster-based ensemble weather forecast analyses. *IEEE Transactions on Visualization and Computer Graphics 24*, 1 (2018), 109–119. doi:10.1109/TVCG.2017.2745178. 3

[Lab18] LABORATORY E. S. S.: European severe weather database, 2018. URL: https://www.eswd.eu/#map_div. 1

[Met22] METEOSCHWEIZ: *Klimabulletin Februar 2022*. 2022. 6

[ML19] MCNABB L., LARAMEE R. S.: Multivariate maps—a glyph-placement algorithm to support multivariate geospatial visualization. *Information 10*, 10 (2019), 302. 3

[NOV*20] NAVARRA C., OPACH T., VROTSOU K., JOLING A., WILK J., NESET T. S.: Visual Exploration of Climate-Related Volunteered Geographic Information. In *Workshop on Visualisation in Environmental Sciences* (2020). doi:10.2312/envirvis.20201092. 2

[NVO*21] NAVARRA C., VROTSOU K., OPACH T., JOLING A., WILK J., NESET T.-S.: A progressive development of a visual analysis interface of climate-related VGI. *Environmental Earth Sciences 80*, 20 (2021), 684. URL: https://doi.org/10.1007/s12665-021-09948-1, doi:10.1007/s12665-021-09948-1. 2

[RBS*18] RAUTENHAUS M., BÖTTINGER M., SIEMENM S., HOFFMAN R., KIRBY R. M., MIRZARGAR M., RÖBER N., WESTERMANN R.: Visualization in meteorology—a survey of techniques and tools for data analysis tasks. *IEEE Transactions on Visualization and Computer Graphics 24*, 12 (December 2018), 3268–3296. doi:10.1109/TVCG.2017.2779501. 2

[RHWS22] RAU S., HEYEN F., WAGNER S., SEDLMAIR M.: Visualization for ai-assisted composing. In *Ismir 2022 Hybrid Conference* (2022). 3

[SMS21] SAHANN R., MÜLLER T., SCHMIDT J.: Histogram binning revisited with a focus on human perception. In *2021 IEEE Visualization Conference (VIS)* (2021), pp. 66–70. doi:10.1109/VIS49827.2021.9623301. 5

[UCAa] UCAR UNIDATA: AWIPS: Advanced Weather Interactive Processing System. https://www.unidata.ucar.edu/software/awips2/. Accessed: 2020-04-28. 2

[UCAb] UCAR UNIDATA: WarnGen Walkthrough - Unidata AWIPS User Manual. http://unidata.github.io/awips2/cave/warngen/. Accessed: 2020-04-30. 2

[WCD*19] WENG D., CHEN R., DENG Z., WU F., CHEN J., WU Y.: Srvis: Towards better spatial integration in ranking visualization. *IEEE Transactions on Visualization and Computer Graphics 25*, 1 (2019), 459–469. doi:10.1109/TVCG.2018.2865126. 3

[Wet] WETTERDIENST D.: Ihre Unwettermeldung. URL: https://www.dwd.de/DE/wetter/warnungen_aktuell/unwettermeldung/unwettermeldung_start_node.html. 1

[WXS12] WENG J., XU Y., SHARMA A. R.: Epidemic analysis and visualization based on digital earth spatio-temporal framework. In *2012 IEEE International Geoscience and Remote Sensing Symposium* (2012), pp. 7220–7223. doi:10.1109/IGARSS.2012.6351996. 3

[ZCL*21] ZHANG M., CHEN L., LI Q., YUAN X., YONG J.: Uncertainty-oriented ensemble data visualization and exploration using variable spatial spreading. *IEEE Transactions on Visualization and Computer Graphics 27*, 2 (2021), 1808–1818. doi:10.1109/TVCG.2020.3030377. 3

[ZKT21] ZHU Y., KANCHARLA P. R., TALLURU C. S. K.: A taxonomy of spatial-temporal data visualization. In *2021 25th International Conference Information Visualisation (IV)* (2021), pp. 223–228. doi:10.1109/IV53921.2021.00043. 3

[ZLG*21] ZHAO J., LIU X., GUO C., QIAN Z. C., CHEN Y. V.: Phoenixmap: An abstract approach to visualize 2d spatial distributions. *IEEE Transactions on Visualization and Computer Graphics 27*, 3 (2021), 2000–2014. doi:10.1109/TVCG.2019.2945960. 3