



City Research Online

City St George's, University of London

Citation: Slingsby, A., Reeve, R. & Harris, C. (2023). Gridded Glyphmaps for Supporting Spatial COVID-19 Modelling. Paper presented at the IEEE VIS: Visualization & Visual Analytics, 22-27 Oct 2023, Melbourne, Australia. doi: 10.1109/VIS54172.2023.00009

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/31115/>

Link to published version: <https://doi.org/10.1109/VIS54172.2023.00009>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Gridded Glyphmaps for Supporting Spatial COVID-19 Modelling

Aidan Slingsby*
Department of Computer Science
City, University of London

Richard Reeve
School of Biodiversity, One
Health & Veterinary Medicine
University of Glasgow

Claire Harris
Biomathematics and Statistics Scotland,
James Hutton Institute




Figure 1: “Infection by age” glyph at a high spatial resolution across a region. Turquoise is ‘susceptible’ (uninfected), pink is ‘dead’, see Fig. 3c for others. Each glyph shows the proportion of the population in each age group (younger at top; older at bottom) in each infection category for days 17, 21 and 25. Temporal lags across the region are apparent as are more deaths for older age groups. Fainter squares have lower population densities, allowing us to see that the disease is spreading slower to smaller populations.

ABSTRACT

We describe our use of gridded glyphmaps to support development of a repurposed COVID-19 infection model during the height of the pandemic. We found that gridded glyphmaps’ ability to interactive summarise multivariate model input, intermediate results and outputs across multiple scales supported our model development tasks in ways that the modellers had not previously seen. We recount our experiences, reflect on potential to support more spatial model development more generally and suggest areas of further work.

Index Terms: Human-centered computing—Visualization—Visualization techniques

1 INTRODUCTION

Spatial models explicitly model spatial processes. As with many types of model, they can help simulate specific scenarios, but they can also be used to help understand spatial processes more generally. They are used in a wide variety of application areas, including climate modelling, population demographics and epidemiology, and they operate at increasingly high spatial resolutions.

During the COVID-19 pandemic, there was a scramble to adapt and develop models capable of helping us understand the nature of the spread, population-level effects and possible effects of interventions [3]. Model developers need to be able to visually explore the multivariate inputs, outputs and intermediate results to help develop a more intuitive understanding of the model dynamics, validate the outputs and choose suitable model parameters.

We (the authors) worked closely together as one of a number of specialist modelling groups within the Scottish COVID-19 Response Consortium (SCRC) [14, 18] which repurposed, adapted and developed new models to help understand the COVID-19 pandemic at a time when little was known about it. The authors’ group

*e-mail: a.slingsby@city.ac.uk

comprised Harris and Reeve adapting and further developing their EcoSYSTEM.jl model [8] – originally for simulating the dynamics of plant species growth, competition and reproduction – working closely with Slingsby to co-design effective interactive analytical visualisation to support model development.

Our input data were demographic data classified by gender and in 10-year age bands, spatially resolved at a 1km² resolution along with the disease parameters and location of initial infection(s). Our output data added infection status (see list in Fig. 3c) over a 30 day window (with a 1-day-timestep). The visual depiction of high-resolution multivariate spatial data presents visualisation challenges, leading to model developers often aggregating spatial areas into coarser (often very coarse) regions, or even producing non-spatial global summaries and univariate summary statistics [6]. Such aggregation hides important spatial and multivariate structure that might be helpful for model validation. The ability to visually explore multivariate spatial summaries of model inputs, outputs and intermediate results became the focus of our work.

Our approach uses “glyphmaps” [20] that can represent spatial variation in multivariate data. They do this by placing glyphs (mini-charts) on maps where each glyph represents a multivariate area-based summary. “Gridded glyphmaps” [15] use a grid-based geographical discretation, creating grid cells that are the bases for the areal summaries represented by glyphs. The addition of zoom/pan interactions enables spatial patterns to be assessed at different scales.

Our contributions are to: (a) make the case for gridded glyphmaps to support development and validation of a COVID-19 model; (b) describe and discuss our visual encodings and interaction designs for helping us achieve this; (c) consider gridded glyphmaps’ wider potential for supporting spatial modelling; and (d) provide open-source implementation for experimentation.

2 TASKS AND REQUIREMENTS

We started working together just as the first version of the model was being created and as modeller needs were emerging. High-level tasks can be summarised in the following:

T1 What is the spatial distribution of the multivariate model




Figure 2: Population glyphs showing comparatively young populations in Edinburgh and Glasgow (young to old from top to bottom).

inputs, (internal) intermediate values and outputs? The inputs and intermediate values drive the model, so understanding these is important for knowing what to expect from the outputs. For example, regions with a higher proportion of older population would expect more severe outcomes.

T2 Which (specific) parameter values are needed for (specific) model processes to have the expected effect? For example, the internal “force of infection” parameter needs to be high enough for COVID-19 infection to spread outside of urban areas, as observed in the case data.

T3 How do the outputs of different models with different scenarios and parameters compare? This can help validate the model itself (e.g. to study effects of adding, removing or re-weighting a specific modelling processes, or exploring parameter spaces) and apply the model operationally; an example of the latter is to assess how variations in lockdown rules affects potential for infection to spread.

3 INTERACTIVE GRIDDED GLYPHMAP APPROACH

Our requirements need a solution that enables multivariate model inputs and outputs to be summarised spatially and to enable this to be done at multiple scales from national (Scotland-wide) all the way down to small towns.

Our solution for **visually summarising spatial multivariate data** is to use a “gridded glyphmap” [7, 15, 20] that regularly grids space, aggregating data within each cell and embedding a multivariate glyph in each. Regular gridding is a good approach when using a high-resolution spatial model. Our solution for **our comparison tasks** was to design comparison glyphs (Fig. 3d). Our solution to **support exploration at multiple spatial scales** is the use of a kind of semantic zoom [2] in which the spatial scale is interactively determined by the zoom level. The fixed gridded discretisation of screen space produces grid cells of fixed screen position and size, but when the underlying data are zoomed and panned, they are reaggregated on-the-fly, revealing data at an appropriate spatial scale for the zoom level, all the way down to the individual 1km² original data values. Thousands of data records are summarised as a single glyph in each cell when zoomed out, whereas at the highest zoom levels a glyph will represent just one record.

3.1 Glyph designs

Designs are based on our tasks (as recommended by Macguire *et al* [10]) and influenced by both Chang *et al*'s [4] design aspects and well-known perceptual guidance on visualisation [11]. Population size/proportion/comparison is the most important quantitative variable and is therefore represented by the most effective visual variable: aligned length. The similarly important ordinal ‘age group’ category is represented by y-position in age order. The 8 nominal ‘infection’ categories are adequately distinguished by hue using ColorBrewer [9] and order in the glyph, which also corresponded to the progression of the disease. Time is represented as animation or on the *x*-axis. We keep the same layout for each grid cell so that




Figure 3: Glyphs used in gridded glyphmaps that supported COVID modelling [5, 16] in two versions, one based on counts and one based on proportions. Opacity indicates population size in the latter, de-emphasising small sample size proportions. (<https://www.staff.city.ac.uk/~sbbb717/glyphmaps/covid/>).

they are learnable [4]. More visually complex glyphs such as Fig. 3c and 3d need larger grid cells to be legible, recognising the important trade-off of spatial precision with legibility.

Population glyphs. Demographics is an important driver of the model and these glyphs were designed to show the structure of the population (Fig. 2) used as *inputs to the model* (a T1 task). Vertical position indicates the age band and horizontal length indicates the size of the sub-population.

Infection by age glyph. The glyphs in Fig. 3b show the *model outputs*, i.e. the number of people in each infection state and by age group for a single day, depicted as stacked barcharts. As for the population glyphs, vertical position encodes the age band, horizontal length indicates size of sub-populations and hue indicates infection state. These animate though each day (see web link in Fig. 3’s caption), conveying lags in spatial spread from seed locations. The population proportion version uses a spine plot that encodes proportions that better indicate the different risks associated with each age group, supporting T1 and T2 tasks.

Infection timeseries glyph. Since animation is not an effective way to help discover trends [13], we designed timeseries-based area charts. To limit glyph complexity we aggregated-out the age-bands. Horizontal position indicates day from 0 (left) to 30 (right), vertical height indicates population and hue indicates the infection state. Fig. 3c shows spatial differences in the rate of infection and its effect on the population over time, supporting T1 and T2.

Infection comparison timeseries glyph. These glyphs encode the *difference between model runs* supporting T3 (comparison) tasks using vertical position above and below the horizontal centre line to indicate whether there were more or fewer people (respectively) in




Figure 4: Relative version of “infection timeseries” glyphs for the most populated area of Scotland. Glyphs representing low populations are faint. Time tick-lines indicating the 9-day position on the x-axes based on mouse pointer position, helps determine time-lags.

the infection state than the model run used for the baseline. Time is on the horizontal axis and hue indicates the infection state. The glyph in Fig. 3d indicates that the two model runs have the same outcome until about day 10, after which that is a larger infection peak (above the horizontal centre-line), with lower ‘susceptible’ (not infected) population (below the horizontal centre-line), though interestingly, deaths are lower (pink below the line).

3.2 Proportions and denominators

A recurring issue with the visual depiction of population-related data is that population size varies between spatial units, particularly for those that are regularly-sized. Expressing counts of ‘infected’ people alone is often not useful where population size varies markedly between spatial units. Expressing as proportions of total population is more informative. However, where overall population sizes are low, such proportions are not statistically stable. Visvalingam [19] suggested using a chi-squared-based statistic to emphasise proportions with larger denominators. We do so in Fig. 3 using transparency to de-emphasise proportions with low population size, separating proportion (size) and denominator (opacity).

3.3 Interactions

Spatial resolution: Semantic zooming interactively changes the spatial resolution of data summarised in the grid cell, whilst keeping the screen-size constant. Analysts can also interactively change the screen-based size of the cells in which there is a trade-off between spatial resolution and size of grid cell therefore complexity of supported glyph [15], an issue that notably applies to the “infection comparison glyph” as reflected later

Scaling: Analysts can rescale the length, colour and/or opacity of glyphs based on the maximum value in view (scaling from zero). This scaling persists with zooming and panning so that magnitudes can be compared until explicitly rescaled. In addition, analysts can interactively increase or decrease this value, allowing values at the lower end of the scale to be explored when required.

Time: Simple mouse movements to scroll through time in the animated glyphs and time tick-marks in the timeseries glyphs, enable

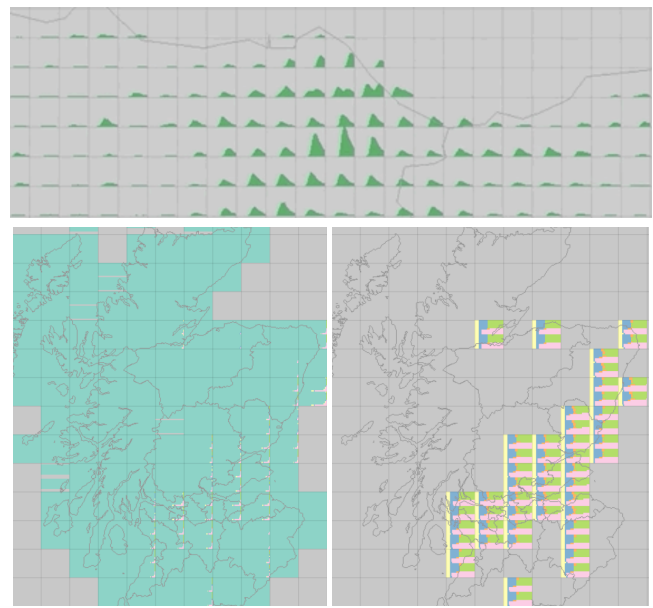


Figure 5: *Top:* Timeseries of the internal “force of infection” value that drives outputs, some of which have double-peaks. *Bottom:* A model run where “force of infection” was too low for infection to spread outside populated areas, contrary to the case data at the time. ‘Susceptible’ has been hidden on the right, reinforcing this observation.

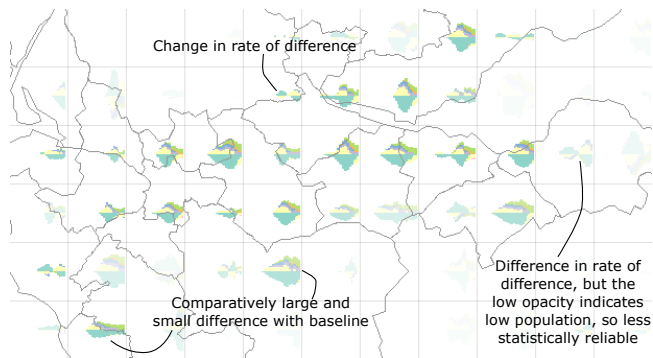


Figure 6: Comparison between between two model runs.

spatial comparison of temporal lags.

Filtering: Infection categories can be switched on and off. ‘Susceptible’ is turned off in Fig. 5 so that the other infection categories can be focused on.

Details on demand: Tooltips indicate data values at the mouse position in any glyph.

4 USAGE, REACTION AND REFLECTION

Reaction and reflection has been published [5, 17]; we focus on usage, reactions and reflections that may inform gridded glyphmaps’ use for other spatial modelling.

4.1 Gridded glyphmaps

Spatial model development usually requires outputs to be separately processed and aggregated to produce static, aggregated and univariate plots [6]. Interactive gridded glyphmaps enabled us to consider richer and more integrated summaries of outputs across multiple spatial scales making it possible to investigate data and model outputs

more freely to better understand how models function, in ways that co-authors Harris and Reeve had not previously seen.

4.2 Tasks and usage

T1: Multivariate spatial distribution. Fig. 4 (“infection timeseries” glyph; time on the x -axis) shows spatial differences in the timeseries of proportional infection. It is a useful overview of temporal patterns across space. The time tick-lines which follow the cursor helped identify time-lags in COVID spread. However, the time animation in Fig. 1’s “infection by age” glyphs – which can be interactively “played” using the mouse to move back – helped identify spatiotemporal dynamics, particularly at the fine spatial resolution in Fig. 1. There is scope for designing glyphs that more specifically focus on small differences in time lags, such as the number of days between peak infection and the average time of peak infection. The animated “infection by age” glyphs also reveal differences in infection by age group – aggregated out of the “infection timeseries” glyphs – with deaths occurring in older populations in Fig. 1. Using these in conjunction with the raw demographic glyphs in Fig. 2 helped us explain spatial differences in how populations responded.

T2: Helping choose appropriate parameter values. In the early stages of model development, we scrutinised the internal intermediate “force of infection” value (Fig. 5; top) that drives model output [17] to help determine mechanism of operation. Through the “infection timeseries” glyphs we were able to see that this parameter needed to be tuned high enough so that COVID-19 infection could ‘escape’ from urban areas into the surrounding countryside (Fig. 5; bottom); hiding the susceptible make this more obvious) as was observed by the case data at the time. We then investigated the inclusion of new data into the model, such as age-specific parameters, age mixing matrices, commuting and pollution data, and the mechanisms by which to incorporate such data. Our “infection by age glyph” was particularly useful for validating differential mixing and death between the different age categories. We found that more mixing between adults produces a corresponding increase in infection amongst these categories, or that an increased probability of death in the older population translated to an increase in those categories that died from COVID-19. Incorporation of further datasets increased the complexity of the model considerably. The inclusion of commuting data and long-distance spread of force of infection decreased our need for tuning this parameter to escape urban areas.

T3: How do different model outputs compare? The “infection comparison timeseries” glyphs show differences between model runs. In particular, we explored the incorporation of environmental variables into COVID-19 dynamics in the model development stage, including weather and pollution. Pollution – especially particulate matter which enters the lungs – is thought to play a role in increased susceptibility and severity of the disease and involved a substantial re-write of the model with new mechanisms. Our comparisons were of runs where we varied a parameter that controlled the impact of pollution on processes like death and hospitalisation. Initial visualisations using the comparison glyphs indicated that these mechanisms were not functioning as intended and flagged this for bug fixing.

4.3 Glyphs and interactions

Developers and users of spatial models tend to focus on highly aggregated univariate summaries [6]. The ability to interact and interrogate the rich multivariate data associated with spatial models is relatively new to spatial modelling. Not compromising on either the time or space dimensions allowed us to include both at varying scales and resolutions. This allowed both the interpretation of ‘big picture’ dynamics and individual small scale variations. The animated glyphs included an option to scroll through time using the mouse to look at the changing course of the epidemic in the simulation. In addition to the elements that inspired ease of use and interpretation of complex data, there were also interactions that aided the user in focusing on

particular aspects of the model inputs and outputs. This included an option to fade data according to their relative population sizes and filtering for specific infection categories. We found this particularly useful during the model comparisons, for which we may want to focus on specific categories according to the parameter being varied, e.g. death, and to concentrate upon differences that are relatively large in comparison to the population size.

5 ISSUES, REFLECTIONS AND FURTHER WORK

Glyph complexity and spatial resolution tradeoff. Larger grid cells help make visually-complex glyphs legible (e.g. the “infection comparison” glyphs) but at the expense of spatial resolution. Interactive techniques for changing the screen resolution (grid size) and changing the geographical resolution (zooming) enables us appropriately trade these off during visual analysis.

Background map. Overlaid administrative boundaries provide visual spatial reference with minimal interference. The subsequent Javascript implementation (below) uses multiscale mapping tiles but adequate gaps between glyphs and/or transparency are necessary. A simple map tile design reduces interference and helps produce visual spatial reference at multiple scales.

Heterogeneity within grid cells. Grid cells may aggregate quite diverse distributions. The high geographical resolution in Fig. 1, reveals nuance that is aggregated out of coarser scale views. Although this is by design, designing glyphs that convey the heterogeneity within cells could alert analysts where to inspect at a finer resolution, for example, showing the data by quartile. Glyphs designs and interactions for investigating this is good further work.

Spatial anomalies. Anomalies might be averaged out by area-based gridding. Designing glyphs that identify buried anomalies would be very useful, if we know how to quantify them. Anomalies often depend on spatial or temporal context, so values that are different from expected based on locally/geographically/temporally weighted statistics or models [1] may be useful anomalies. For example, we might be interested unusually high death rates for age groups given the local average or what happened in previous days.

Modifiable Areal Unit Problem (MAUP). Imposed regular gridding is that arbitrary and does not respect geographical context, potentially splitting important populations and leading to the Modifiable Areal Unit Problem [12]. The repeated reaggregation of gridded glyphmaps makes them particularly vulnerable to this. It is particularly problematic when there is an arbitrariness about the aggregation (the case with regular gridding) and that aggregation is retained. Glyph stability when interactively panning gives a visual indication of the impact of MAUP. There is scope for further work in quantifying this and depicting the effects through glyphs.

Model mechanisms. Spatial models model spatial processes. They may be particular ‘tipping points’ that govern which process dominates. There is scope future work to design glyphs that summarise the spatial distribution of details of model mechanisms.

Implementation. The original Java implementation is available at <https://www.staff.city.ac.uk/~sbbb717/glyphmaps/covid/>. We subsequently produced an open-source Javascript implementation within an Observable notebook environment (<https://observablehq.com/collection/@aidans/covid-19-modelling>). This can be to be explored and experimented with for these and other spatial datasets.

6 CONCLUSION

We demonstrate a case for interactive gridded glyphmaps for supporting the development of spatial models through our COVID-19 modelling case study. We have discussed visual coding and interaction designs that supported our model development tasks. This demonstrates the utility gridded glyphmaps for the development of spatial model more generally. Our implementations can be used to experimentation and we have suggested lines of future research.

ACKNOWLEDGMENTS

This work was undertaken in part as a contribution to the Rapid Assistance in Modelling the Pandemic (RAMP) initiative, coordinated by the Royal Society, and we gratefully acknowledge all involved in the Scottish COVID-19 Response Consortium that formed during RAMP. We were also supported by funding from the Engineering and Physical Science Research Council (EP/V054236/1 and EP/V033670/1) and STFC (ST/V006126/1).

REFERENCES

- [1] C. Brunson, S. Fotheringham, and M. Charlton. Geographically weighted regression. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 47(3):431–443, 1998.
- [2] T. Buering, J. Gerken, and H. Reiterer. User interaction with scatterplots on small screens - a comparative evaluation of geometric-semantic zoom and fisheye distortion. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):829–836, 2006. doi: 10.1109/TVCG.2006.187
- [3] M. Chen, A. Abdul-Rahman, D. Archambault, J. Dykes, A. Slingsby, P. D. Ritsos, T. Torsney-Weir, C. Turkay, B. Bach, A. Brett, H. Fang, R. Jianu, S. Khan, R. S. Laramée, P. H. Nguyen, R. Reeve, J. C. Roberts, F. Vidal, Q. Wang, J. Wood, and K. Xu. Rampvis: Towards a new methodology for developing visualisation capabilities for large-scale emergency responses, 2020.
- [4] D. H. Chung, P. A. Legg, M. L. Parry, R. Bown, I. W. Griffiths, R. S. Laramée, and M. Chen. Glyph sorting: Interactive visualization for multi-dimensional data. *Information Visualization*, 14(1):76–90, 2015.
- [5] J. Dykes, A. Abdul-Rahman, D. Archambault, B. Bach, R. Borgo, M. Chen, J. Enright, H. Fang, E. E. Firat, E. Freeman, T. Gonen, C. Harris, R. Jianu, N. John, S. Khan, A. Lahiff, R. Laramée, L. Matthews, S. Mohr, P. H. Nguyen, A. A. Rahat, R. Reeve, P. D. Ritsos, J. C. Roberts, A. Slingsby, B. Swallow, T. Torsney-Weir, C. Turkay, R. Turner, F. P. Vidal, Q. Wang, J. Wood, and K. Xu. Visualization for epidemiological modelling: challenges, solutions, reflections recommendations. *Philosophical Transactions of the Royal Society A*, 3 2022.
- [6] S. Grainger, F. Mao, and W. Buytaert. Environmental Modelling Software Environmental data visualisation for non-scientific contexts : Literature review and design framework. *Environmental Modelling and Software*, 85:299–318, 2016. doi: 10.1016/j.envsoft.2016.09.004
- [7] G. Grolemond and H. Wickham. Visualizing complex data with embedded plots. *Journal of Computational and Graphical Statistics*, 24(1):26–43, 2015.
- [8] C. Harris, R. Reeve, J. Cook, S. Lovett, E. Perim Martins, A. Robson, and B. Wu. EcoSYSTEM.jl. <https://github.com/ScottishCovidResponse/EcoSYSTEM.jl>, 2020.
- [9] M. Harrower and C. A. Brewer. Colorbrewer.org: An online tool for selecting colour schemes for maps. *The Cartographic Journal*, 40(1):27–37, June 2003. doi: 10.1179/000870403235002042
- [10] E. Maguire, P. Rocca-Serra, S.-A. Sansone, J. Davies, and M. Chen. Taxonomy-based glyph design—with a case study on visualizing workflows of biological experiments. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2603–2612, 2012.
- [11] T. Munzner. *Visualization analysis and design*. CRC press, 2014.
- [12] S. Openshaw. The modifiable areal unit problem. *Quantitative geography: A British view*, pp. 60–69, 1981.
- [13] G. G. Robertson, S. K. Card, and J. D. Mackinlay. Information visualization using 3d interactive animation. *Commun. ACM*, 36(4):57–71, apr 1993. doi: 10.1145/255950.153577
- [14] SCRC. The scottish covid-19 response consortium. <https://www.gla.ac.uk/research/az/scrc/>, 2020.
- [15] A. Slingsby. Tilemaps for summarising multivariate geographical variation. VISREG (Visual Summarization and Report Generation: Beyond Scatter-Plots and Bar-Charts) Workshop at IEEE VIS, 2018.
- [16] A. Slingsby, C. Harris, and R. Reeve. Demo: Glyphmaps for population and covid. <https://www.staff.city.ac.uk/~sbbb717/glyphmaps/covid/>.
- [17] A. Slingsby, C. Harris, and R. Reeve. Rampvis: Gridded glyphmaps. <https://observablehq.com/@aidans/rampvis-idiom-gridded-glyphmaps>.
- [18] The Royal Society. Rapid assistance in modelling the pandemic: Ramp. <https://epcced.github.io/ramp/>, 2020.
- [19] M. Visvalingam. The signed chi-square measure for mapping. *The Cartographic Journal*, 15(2):93–98, 1978. doi: 10.1179/caj.1978.15.2.93
- [20] H. Wickham, H. Hofmann, C. Wickham, and D. Cook. Glyph-maps for visually exploring temporal patterns in climate data and models. *Environmetrics*, 23(5):382–393, 2012.