Intelligent Geovisualizations for Open Government Data (Vision Paper)

Auriol Degbelo

Institute for Geoinformatics, University of Münster
degbelo@uni-muenster.de

Christian Kray

Institute for Geoinformatics, University of Münster
c.kray@uni-muenster.de

ABSTRACT

Open government datasets (OGD) have been flooding the Web in recent years. Geovisualisations are the natural way of making sense of them, and have been gradually coming out. However, one key problem is the lack of flexibility of these visualizations, which severely limits their re-use in new scenarios. This article therefore proposes to *increase the intelligence* of existing geovisualisations by incorporating five features, to make better use of OGD: (i) automatic geographic data type recognition, (ii) generation of geovisualisation designs, (iii) monitoring of users' understanding of geographic facts, (iv) self-optimization, and (v) user activity recognition. In addition to benefiting users of OGD, realizing these features presents rich scientific challenges and opportunities for Geovisualization research, the OGD landscape (and beyond).

CCS CONCEPTS

• **Human-centered computing** → **Geographic visualization**; **Visualization systems and tools**;

KEYWORDS

Geovisualisation, Intelligent Systems, Adaptive Systems, Computer-Generated User Interfaces, Open Government Data

ACM Reference Format:

Auriol Degbelo and Christian Kray. 2018. Intelligent Geovisualizations for Open Government Data (Vision Paper). In *26th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (SIGSPATIAL '18)*, November 6–9, 2018, Seattle, WA, USA. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3274895.3274940>

1 MOTIVATION AND BACKGROUND

Open data laws, political movements and other drivers have led to the increasing availability of public and governmental data. Open Data Inception (opendatainception.io) lists no less than 2600 open data portals around the world, covering a broad range of topics such as education, weather, population, environment and health, to name just a few. Some of these portals (e.g., the European Open Data Portal, <https://www.europeandataportal.eu/> or the OpenDataSoft's data network, <https://data.opendatasoft.com/>) provide widgets to visualize these datasets. Since in many cases, data is linked to spatial concepts (such as zip codes, districts or even coordinates) the

relevancy of geovisualizations in this context has increased as well. This trend has, however, also highlighted some key issues in this context. In particular, the (geo-)visualizations for different datasets vary, and if untrained users try to generate them from raw datasets there are few if any safeguards to prevent the selection of inappropriate visualizations. These issues make it difficult for citizens to understand the meaning of, and to compare different datasets (e.g., from the same governmental body or from different ones).

There are at least three reasons why more versatile (or flexible) visualizations will have a positive impact on the Open Government Data (OGD) landscape: 1) there is a need (confirmed in a previous survey, see [17]) of a portion of the OGD consumers' population not just to consume existing visualizations, but to bring in their own data and visualize it as they like; 2) the expertise to produce a visualization from a dataset is not always available; and 3) even if the expertise (e.g., programming skills) is there, the time to create a new visualization is not always available.

This article presents a vision of intelligent geovisualizations for open government data in order to inspire future research that will make OGD more useful and accessible to a broad audience. The vision covers five main aspects: (i) automatic geographic data type recognition, (ii) generation of geovisualization designs, (iii) monitoring of users' understanding of geographic facts, (iv) self-optimization, and (v) user activity recognition. Intelligence is defined after Albus [1] as “the ability of a system to act appropriately in an *uncertain* environment” (emphasis added). For each aspect, we first introduce the ‘uncertainties’ the geovisualization needs to cope with, followed by existing work on which future endeavours can build upon to realize intelligent OGD geovisualizations.

We use the term ‘Geovisualization’ in line with [28] to broadly include interactive maps, network graphs, charts/graphs, tables, symbols, diagrams and pictures. Geovisualizations can act as catalysts for citizen engagement in the OGD landscape [15]. In keeping with Albus [1], intelligence requires at least the ability to sense the environment, make decisions, and control action. The sort of intelligence to be attained in machines does not have to simulate human intelligence (see [36]). ‘Open Government Data’ is an inherently ambiguous term (see e.g., [43]), but is used here to denote public sector data which is freely available for re-use. Following Roth [30], a user (i.e., an open government data consumer) is modelled as having three characteristics: expertise, ability, and motivation. The anticipated main beneficiary of this vision has low expertise in programming and geovisualization, possibly low/medium spatial thinking abilities, and desires to consume open geographic data out of curiosity. Though the features of an intelligent geovisualization are all introduced at once here (see Figure 1), it is not required that a single geovisualization has them all. Arguably, any geovisualization which supports any of the features presented next, exhibits

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGSPATIAL '18, November 6–9, 2018, Seattle, WA, USA

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-5889-7/18/11.

<https://doi.org/10.1145/3274895.3274940>

Figure 1: From open government data to intelligent geovisualizations. Dotted arrows suggest implicit or explicit input from the user in the process of realizing the feature.

some sort of intelligence. The more features supported, the more intelligent the geovisualization.

2 RECOGNIZE GEOGRAPHIC DATA TYPES

If an untrained user is to be able to easily get a useful visualization for an open geographic dataset, a critical feature of the next generation of geovisualizations would be the ability to automatically recognize the *semantic type* of the geodataset at hand. This task is fraught with much uncertainty for the computer (and the untrained data analyst) because there are diverse semantic types for spatial datasets. Ferreira et al. [16] identified time series, trajectories and coverage as basic types of observations; Scheider et al. [32] proposed fields, inverted fields, lattices, events, and objects; Stevens [35]’s distinction between nominal, ordinal, interval and ratio data is also relevant in the context of spatial information. The uncertainties are further aggravated when one considers that there is a many-to-many relationship between semantic (e.g., trajectory) and syntactic types (e.g., GeoJSON, Comma Separated Value, or Resource Description Framework) of geodatasets. We have recently begun experimenting with automatic recognition of Stevens’s four types of datasets, taking a GeoJSON dataset as a starting point. One lesson learned is that automatic semantic geodata type recognition is a human-computation problem (in the sense of [40]). Hence, the approach of algorithmically-guided user interaction proposed recently in [38] is one promising way of approaching this issue.

3 PROPOSE GEOVISUALIZATION DESIGNS

‘Insight is formed through interaction’ [28], and untrained users are interested mostly in the information/insight they can extract from a dataset (not in the raw dataset per se). That is, after having identified the type of spatial data at hand, an OGD geovisualization system should propose meaningful visualization designs pertaining to the current dataset. Here also, there are uncertainties, not so much on the computer side, but more on the side of the untrained data analyst. Indeed, pitfalls and best practices of geovisualizations may be well-known to experts, but unknown to non-experts. For instance, novice map makers have the tendency to overload a single map with too much detail, while it is usually better to err on the side of simplicity and to produce two or three maps, each focused on a single topic [19]. Count data (e.g., number of births in a country) is better represented through a ‘bigger-means-more’ coding in a dot-array map; representing count data using the ‘darker-means-more’ rule in a choropleth map is misleading [24]. In his seminal

work, Bertin [7] suggested eight ‘visual variables’ (i.e., graphical dimensions across which a map/visualization can be varied to encode information) which could be useful while producing automatic designs. Rules prescribing the use of a visual variable given Stevens [35]’s data types were recently summarized in [31]: for instance, color, hue, orientation and shape are useful for encoding nominal information; ordinal information is better encoded using other visual variables such as color value, crispness and transparency.

While these heuristics have been documented in previous work, a consolidated set of guidelines for the design of interactive maps is yet to be produced (see [30]), and the same holds true for other types of geovisualizations. Automatically suggesting designs for datasets would need a democratization of the heuristics listed above, and of further best practices of geovisualizations. Bresciani and Eppler [9] lists common pitfalls of information visualization in general. Their list (and in particular the designer-induced cognitive pitfalls) can be translated into constraints for the automation process so that no design encodes input data in a way that gives rise to a ‘pitfall’.

Rawgraphs [23] provides useful features for the automatic visualization of spreadsheet data, and could be modified for the case of geographic data. The SIGSPATIAL community has done some work on automatic generation of road networks from GPS traces (see e.g., [10, 13]), which addresses one specific type of geovisualization (i.e., network), and might inform future approaches that automatically produce more complex geovisualizations (e.g., interactive maps and timelines). Zavala-Romero et al.’s work on generating web GIS without programming knowledge [44] is in line with the idea outlined here, but only implements a small subset of it (i.e., their tool automatically builds web GIS interfaces to visualize NetCDF data). Existing work on map labelling (e.g., [5, 27]), or computing aggregation of large point datasets [6] is also relevant in this context. Insights from ongoing work on meaningful spatial prediction and aggregation [34] would also help tackle the issue presented here.

4 MONITOR UNDERSTANDING

Once users have a geovisualization design (or several geovisualization designs) for their dataset, the next big question is whether they *truly* understand it. If we are to realize Shneiderman’s vision of designing computer-based tools which ‘enable more people to be more creative more often’ [33], this question is of great importance. Here the uncertainties, primarily on the side of the computer, lie in getting it to determine whether the geovisualization is understandable for the current user. In a classical geovisualization maker/user setting, the geovisualization maker can informally ask the user (or observe her) to assess the extent to which her goal (i.e., transmit geographic insight) is attained (or not). Getting to the point where the geovisualization (reliably) collects this feedback by itself will require years of sustained research effort.

One approach to tackle this issue is the use of computer-generated questionnaires. Using these questionnaires to monitor understandability could drive an ‘anticipation feedback loop’ [18], where the computer modifies and adapts designs until the users’ interpretation corresponds to the one by the computer system. The questionnaire-based insight measurement approach proposed in [11], and the idea of automatically generating insight-questions from the source dataset(s) of a geovisualization, could also be beneficial at this point.

The underlying assumption behind this idea is that a source dataset and its geovisualization can be considered informationally equivalent in the sense of [21]. In terms of ‘metaphors we compute by’ [39], computer-generated questionnaires approach the monitoring issue by gathering *information about users’ understanding from their self-reports*. Another, equally relevant metaphor worth exploring is that of *information about users from behavioral observations*, i.e., explore the extent to which bread crumbs of their requests, keyboard inputs, mouse clicks or other relevant user activities can be used to make reliable inferences about their current level of understanding.

5 OPTIMIZE GEOVISUALIZATION DESIGNS

Once users have provided feedback on how useful the geovisualization is when trying to understand the data at hand, there is the question whether there could be a geovisualization doing an even better job at supporting them. Since they are not aware of all the intricacies of the geovisualization generated, the onus of answering this question is on the computer. That is, the uncertainties arise here mainly on the side of the computer.

Insights from work on computational user interface design (for a review of approaches, see [26]) could be useful while addressing this issue. The main question here is that of the objective function for the optimization. Here, the suggestion is that *understanding* on the side of the user is used as objective function (so that the geovisualization ultimately enables people to effectively act intelligently). In the context of Open Government Data, understanding as an objective function is critical: unless there is an (understandable) visualization, knowledge is not really put into the hands of the data consumer (i.e., one of the main goals of opening up government data is *not yet achieved*). Empirical guidelines on factors which facilitate the generation of geographic insight during interaction are scarce [30], and more work is needed in this area to (i) identify these factors, and (ii) produce algorithms which maximize understanding as an objective function of computer-generated geovisualizations.

Constraint-based visualization systems, which consider both expressiveness and perceptual effectiveness while automatically generating visualizations, are relevant in this context. Expressiveness here refers to the ability of a visualization to convey all facts in the data, and only facts in the data, while effectiveness denotes the visualization’s ability to convey facts in a way that they are readily perceivable by the end users (see [22]). Existing systems (e.g., Voyager [41] or Draco [25]) have been used to produce bar/line charts and scatterplots, and more work will be needed in the future to expand them to more complex interactive geovisualizations.

6 RECOGNIZE USER ACTIVITIES

A fifth, useful feature of the next generation of geovisualizations is to automatically determine what users are trying to do, and support them well in doing that. For example (and as suggested in [2]), multiple different presentations of the same data should be proposed during the data exploration phase. Shortcuts to save (or share) the geovisualization could be hidden by default (reducing the number of symbols to render on the interface), and displayed when it makes most sense (i.e., towards the end of the interaction session). If basic interaction scenarios for a given geovisualization are specified (like in [3]), intelligent guidance could try to fit users’ activities to one

of the most plausible scenario, and use that knowledge to suggest functionalities to try out next at key stages of the users’ navigation. The uncertainties here (mostly on the side of the computer) lie in pinpointing exactly what users are trying to achieve. This is an underexplored area in geovisualization research, and Fabrikant’s early call to ‘formulate domain independent visualization tasks that are generic enough to be effectively shared amongst a heterogeneous user community’ [14] is as topical as it was 20 years ago.

Taxonomies of interaction tasks have begun to emerge, and will be useful in addressing this issue. Roth [29] suggested three broad interaction goals (i.e., procure, predict and prescribe) and five more specific objectives (i.e., identify, compare, rank, associate, and delineate) for geovisualizations based on an empirical user study. Roth’s taxonomy proved usable while characterizing web maps, but still needs refinement so that answering the question of ‘which interaction operators matter for which interaction objective’ becomes possible (see [12]). Kiefer et al. [20] considered six main map-based activities during their work: free exploration, (global) search, route planning, focused search, line following, and polygon comparison. Though they did not provide the rationale for choosing only these six types of activities, their study using eye movement characteristics to automatically detect map users’ activities ended with an overall accuracy of 78%, providing evidence for the pertinence of these six types. Building upon the works mentioned here will be crucial for the implementation of algorithms, which enable automatic user activity recognition for intelligent geovisualizations.

7 FURTHER REMARKS

Besides the five kinds of uncertainties listed above, a number of further features are desirable for the success of the vision. Notably, the geovisualization should tell about its decision tree (and allow users to modify it) - or put differently comply with the principles of ‘algorithmic transparency’ [8]. Algorithmic transparency (which very likely would require visualizations so that untrained professionals can benefit from it) would be helpful to build trust between users and the geovisualization - and indirectly between users and the (public) institution (or community) using it as a means to make its data more understandable. Furthermore, multi-device portability (e.g., building on the early idea of interface plasticity, see [37]) and multi-modality support (e.g., account for speech and hand gesture) could be convenient add-ons, though not discussed explicitly here.

Finally it’s worth mentioning few key differences between this vision, and the occurrence of the term in previous articles. Yingjie et al. [42] proposed a related discussion on adaptive geovisualizations, focusing mainly on architectural aspects. Instead, the focus of this work is deliberately functional, and the paper has brought forth five *user requirements* of intelligent geovisualizations. Using the verb form to formulate the title of each of the previous sections is intentional, and helps to highlight this (‘the system should recognize the geographic type of my current dataset’, ‘propose (to me) meaningful geovisualization designs’, and so on). Furthermore, the article used the existing literature to point out where we currently stand, and to outline ways forward (highlighting that much still needs to be done to get where we want to be). Andrienko and Andrienko [4] proposed a system which can be viewed as an early implementation of the ‘generation of geovisualization designs’ feature

mentioned above. Their distinction between domain-dependent and domain-independent components of an intelligent visualization is useful, and this article has outlined steps currently needed to further advance research on the latter type of component, pointing out that recent progress in geo-ontologies, information visualization research, computational user interface design, and the science of cartographic interaction offer new prospects for tackling the issue of automatically picking the right geovisualization for a given task.

8 CONCLUSION

Open Government Data (OGD) has been flooding the Web in recent years, and geovisualizations are gradually coming out to help make sense of them. This article has suggested increasing the intelligence of these geovisualizations to enhance their flexibility (and ability to be re-used in several scenarios). The vision covers five main aspects: automatic geographic data type recognition, generation of geovisualization designs, monitoring of users' understanding of geographic facts, self-optimization, and user activity recognition. In modelling intelligence as the ability to act appropriately in an uncertain environment, the work provides a clear criterion to distinguish non-intelligent geovisualizations from intelligent ones (or less intelligent geovisualizations from more intelligent ones). The vision has taken OGD as main scenario, but intelligent geovisualizations would be beneficial in a number of other areas, notably data journalism (recognize spatial data types & propose geovisualization designs), E-participation (monitor understanding), computer-assisted spatial learning (optimise geovisualization designs) and pedestrian/car navigation (recognize user activities). In sum, intelligent geovisualizations are essential to make OGD understandable and offer many opportunities for future research.

Acknowledgements: Funding from the European Union through the GEO-C project (H2020-MSCA-ITN-2014, Grant Agreement Number 642332, <http://www.geo-c.eu/>) is gratefully acknowledged. We also thank the anonymous reviewers for their feedback.

REFERENCES

- [1] James S Albus. 1991. Outline for a theory of intelligence. *IEEE Trans. on Systems, Man, and Cybernetics* 21, 3 (1991), 473–509. <https://doi.org/10.1109/21.97471>
- [2] Gennady Andrienko and Natalia Andrienko. 1998. Intelligent visualization and dynamic manipulation: two complementary instruments to support data exploration with GIS. In *Proceedings of AVI'98*. 66–75. <https://doi.org/10.1145/948496.948506>
- [3] Natalia Andrienko and Gennady Andrienko. 2001. Intelligent support for geographic data analysis and decision making in the web. *Journal of Geographic Information and Decision Analysis* 5, 2 (2001), 115–128.
- [4] Natalia Andrienko and Gennady Andrienko. 2007. Intelligent visualisation and information presentation for civil crisis management. *Transactions in GIS* 11, 6 (dec 2007), 889–909. <https://doi.org/10.1111/j.1467-9671.2007.01078.x>
- [5] Lukas Barth, Benjamin Niedermann, Martin Nöllenburg, and Darren Strash. 2016. Temporal map labeling: a new unified framework with experiments. In *Proceedings of SIGSPATIAL 2016*. 1–10. <https://doi.org/10.1145/2996913.2996957>
- [6] Christian Beilschmidt, Thomas Fober, Michael Mattig, and Bernhard Seeger. 2017. A linear-time algorithm for the aggregation and visualization of big spatial point data. In *SIGSPATIAL 2017*. 73:1–73:4. <https://doi.org/10.1145/3139958.3140037>
- [7] Jacques Bertin. 1983. *Semiology of graphics: diagrams, networks, maps (Translated by William J. Berg)*. The University of Wisconsin Press, Madison.
- [8] Robert Brauneis and Ellen P. Goodman. 2018. Algorithmic transparency for the smart city. *Yale Journal of Law & Technology* (2018).
- [9] Sabrina Bresciani and Martin J. Eppler. 2015. The pitfalls of visual representations. *SAGE Open* 5, 4 (dec 2015), 1–14. <https://doi.org/10.1177/2158244015611451>
- [10] Lili Cao and John Krumm. 2009. From GPS traces to a routable road map. In *Proceedings of SIGSPATIAL 2009*. 3–12. <https://doi.org/10.1145/1653771.1653776>
- [11] Auriol Degbelo. 2017. Linked data and visualization: two sides of the transparency coin. In *Proceedings of UrbanGIS'17*. 1–8. <https://doi.org/10.1145/3152178.3152191>
- [12] Auriol Degbelo and Tomi Kauppinen. 2018. Increasing transparency through web maps. In *Companion of Proceedings of the Web Conference 2018*. 899–904.
- [13] David Duran, Vera Sacristán, and Rodrigo Silveira. 2016. Map construction algorithms: an evaluation through hiking data. In *Proceedings of MobiGIS'16*. Burlingame, California, USA, 74–83. <https://doi.org/10.1145/3004725.3004734>
- [14] Sara I Fabrikant. 2001. Building task-ontologies for geovisualization. *ICA Pre-Conference Workshop on Geovisualization on the Web* (2001).
- [15] Thore Fechner and Christian Kray. 2014. Georeferenced open data and augmented interactive geo-visualizations as catalysts for citizen engagement. *eJournal of eDemocracy and Open Government* 6, 1 (2014), 14–35.
- [16] Karine Reis Ferreira, Gilberto Camara, and António Miguel Vieira Monteiro. 2014. An algebra for spatiotemporal data: from observations to events. *Transactions in GIS* 18, 2 (apr 2014), 253–269. <https://doi.org/10.1111/tgis.12030>
- [17] Alvaro Graves and James Hendler. 2013. Visualization tools for open government data. In *14th Annual International Conference on Digital Government Research*.
- [18] Anthony Jameson and Wolfgang Wahlster. 1982. User modelling in anaphora generation: ellipsis and definite description. In *Proceedings of the 5th European Conference on Artificial Intelligence*. Paris, France, 222–227.
- [19] Robert B. Kent and Richard E. Klosterman. 2000. GIS and mapping. *Journal of the American Planning Association* 66, 2 (jun 2000), 189–198.
- [20] Peter Kiefer, Ioannis Giannopoulos, and Martin Raubal. 2013. Using eye movements to recognize activities on cartographic maps. In *Proceedings of ACM SIGSPATIAL 2013*. 488–491. <https://doi.org/10.1145/2525314.2525467>
- [21] Jill H. Larkin and Herbert A. Simon. 1987. Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science* 11, 1 (jan 1987), 65–100.
- [22] Jock Mackinlay. 1986. Automating the design of graphical presentations of relational information. *ACM Transactions on Graphics* 5, 2 (apr 1986), 110–141.
- [23] Michele Mauri, Tommaso Elli, Giorgio Caviglia, Giorgio Ubaldi, and Matteo Azzi. 2017. RAWGraphs: a visualisation platform to create open outputs. In *Proceedings of CHIItaly '17*. Cagliari, Italy, 1–5. <https://doi.org/10.1145/3125571.3125585>
- [24] Mark Monmonier. 2005. Lying with maps. *Statist. Sci.* 20, 3 (2005), 215–222.
- [25] Dominik Moritz, Chenglong Wang, Greg L. Nelson, Halden Lin, Adam M. Smith, Bill Howe, and Jeffrey Heer. 2018. Formalizing visualization design knowledge as constraints: actionable and extensible models in Draco. *IEEE TVCG* (2018).
- [26] Antti Oulasvirta. 2017. User interface design with combinatorial optimization. *Computer* 50, 1 (jan 2017), 40–47. <https://doi.org/10.1109/MC.2017.6>
- [27] Shangfu Peng, Hanan Samet, and Marco Adelfio. 2014. Viewing streaming spatially-referenced data at interactive rates. In *ACM SIGSPATIAL 2014*. 409–412.
- [28] Jonathan Roberts. 2008. Coordinated multiple views for exploratory geovisualization. In *Geographic Visualization*. John Wiley & Sons, Ltd, 25–48.
- [29] Robert E. Roth. 2013. An empirically-derived taxonomy of interaction primitives for interactive cartography and geovisualization. *IEEE Trans. on Vis. and Computer Graphics* 19, 12 (2013), 2356–2365. <https://doi.org/10.1109/TVCG.2013.130>
- [30] Robert E Roth. 2013. Interactive maps: What we know and what we need to know. *JOSIS* 6 (2013), 59–115. <https://doi.org/10.5311/JOSIS.2013.6.105>
- [31] Robert E. Roth. 2017. Visual variables. In *International Encyclopedia of Geography: People, the Earth, Environment and Technology*. John Wiley & Sons, Ltd, 1–11.
- [32] Simon Scheider, Benedikt Gräler, Edzer Pebesma, and Christoph Stasch. 2016. Modeling spatiotemporal information generation. *Int. J. Geogr. Inf. Sci.* 30, 10 (mar 2016), 1980–2008. <https://doi.org/10.1080/13658816.2016.1151520>
- [33] Ben Shneiderman. 2007. Creativity support tools: accelerating discovery and innovation. *Commun. ACM* 50, 12 (dec 2007), 20–32.
- [34] Christoph Stasch, Simon Scheider, Edzer Pebesma, and Werner Kuhn. 2014. Meaningful spatial prediction and aggregation. *Environmental Modelling & Software* 51 (2014), 149–165. <https://doi.org/10.1016/j.envsoft.2013.09.006>
- [35] S S Stevens. 1946. On the theory of scales of measurement. *Science* 103, 2684 (1946), 677–680.
- [36] Alex S. Taylor. 2009. Machine intelligence. In *Proceedings of the 27th international conference on Human factors in computing systems - CHI 09*. 2109–2118.
- [37] David Thevenin and Joëlle Coutaz. 1999. Plasticity of user interfaces: framework and research agenda. In *Proceedings of INTERACT '99*. 110–117.
- [38] Thomas van Dijk and Alexander Wolff. 2017. Algorithmically-guided user interaction. In *Proceedings of ACM SIGSPATIAL 2017 (SIGSPATIAL '17)*. 11:1–11:4.
- [39] Alvaro Videla. 2017. Metaphors we compute by. *Commun. ACM* 60, 10 (2017).
- [40] Luis von Ahn. 2009. Human computation. In *Proceedings of the 46th Annual Design Automation Conference on DAC - DAC '09*. 418–419.
- [41] Kanit Wongsuphasawat, Dominik Moritz, Anushka Anand, Jock Mackinlay, Bill Howe, and Jeffrey Heer. 2016. Voyager: exploratory analysis via faceted browsing of visualization recommendations. *IEEE TVCG* 22, 1 (jan 2016), 649–658.
- [42] Wang Yingjie, Liu Yue, Chen Xiaogang, Chen Yufen, and Liqiu Meng. 2001. Adaptive geovisualization: an approach towards the design of intelligent geovisualization systems. *Journal of Geographical Sciences* 11, S1 (dec 2001), 1–8.
- [43] Harlan Yu and David G. Robinson. 2012. The new ambiguity of 'open government'. *UCLA Law Review* 59, 178 (2012).
- [44] Olmo Zavala-Romero, Eric P. Chassignet, Jorge Zavala-Hidalgo, Panagiotis Velisariou, Harshul Pandav, and Anke Meyer-Baese. 2014. OWGIS 2.0: open source Java application that builds web GIS interfaces for desktop and mobile devices. In *ACM SIGSPATIAL 2014*. 311–320. <https://doi.org/10.1145/2666310.2666381>