

DOI:10.1145/2756546

Digital maps can be engineered to adapt to a person's unique interests and experience in geographic space.

BY ANDREA BALLATORE AND MICHELA BERTOLOTTO

Personalizing Maps

GEOGRAPHIC MAPS CONSTITUTE a ubiquitous medium through which we understand, construct, and navigate our natural and built surroundings. At the intersection of the explosion of geographic information online, data-mining techniques, and the increasing popularity of Web maps, a novel possibility has emerged: Instead of generating one map for large numbers of users, user profiling and implicit feedback analysis can support creation of a different map for each person. The automated personalization of the map-making process is still in its infancy but has the potential to provide more relevant maps to millions of users worldwide.

While mapmaking has traditionally aimed to produce static maps to be printed and distributed to a target audience, geographic information systems (GISs) provide interactive tools to collect and process information dynamically, transforming not only cartography but also geography, urban planning, and any activity that relies on geographic knowledge. Since the 1960s, using GISs, geographers, urban planners, army generals, and economists have been generating different representations of the same input

data to better understand diverse geospatial phenomena. Over the past decade, GISs have further merged with Web technologies and mobile computing, enabling mass adoption of digital maps while overcoming the limitations of paper maps.

As interactive digital maps replace paper maps, this “ubiquitous cartography” is quietly becoming part of our lives, changing not only the consumption but also the production of geographic information.^{5,7} Just below the surface of this tumultuous reconfiguration, the fundamental problems of cartography have hardly changed. Complex, dynamic, and uncertain geographic data needs to be represented on a screen, selecting what needs to be displayed and how to display it with respect to the user's informational needs.⁶ Appropriate cartographic projections, scales, generalization principles, human-computer interaction, and semiotic conventions constitute essential ingredients for the design of usable digital maps.

Although maps are often perceived as a form of objective, scientific knowledge about the world, the same area can be represented from many alternative perspectives, including and ex-

>> key insights

- **Cartography traditionally focuses on producing maps for large groups of readers, and digital maps, including Google's, have barely begun to challenge this approach; collecting explicit and implicit feedback from users, digital cartography is able to capture a person's geographic knowledge, experiences, and attitudes, better supporting spatial learning and decision making.**
- **As a research frontier, automated map personalization requires real-time task detection, geographic user profiling, trajectory analytics, data fusion, geo-visualization, and sentiment analysis, along with insight from cognitive psychology and human geography.**
- **Depending on design, personal maps could foster exploration of the environment beyond the user's known territory or reinforce segregation, fragmenting the collective knowledge of the spaces we inhabit.**



cluding different pieces of uncertain information, and choosing arbitrary graphical and symbolic conventions; for example, Figure 1 includes alternative representations of University College Dublin. For cartographers, it is uncontroversial that radically different maps are needed to perform different tasks. Nautical charts, tourist maps, and urban-planning maps display different geographic information tailored to specific tasks (such as reaching a port safely, understanding the structure of a city, or identifying a suitable location to build a new bridge). Less obvious is the fact that different maps might be needed by different people to perform the same task. Since the 1950s, psychological studies have shown every person perceives and develops an individual mental model of their environment, based on direct and mediated subjective experiences.¹¹

Likewise, since the late 1990s, the economic value of personalization of Web-based services has attracted considerable attention, resulting in now-ubiquitous personalized news stories, commercial offers, film recommendations, and search results. In 1995, Nicholas Negroponte, founder of MIT's Media Lab, imagined a newspaper called *Daily Me* that would automatically collect and arrange stories relevant to the reader, rather than impose the same content on everyone, overcoming the paradigm of mass production that dominated the 20th century.¹⁶ Knowing a customer's behavior and tastes through surveillance techniques has become commonplace in marketing, in a tight feedback loop between companies and their current or prospective consumers, in what has been called by the oxymoron "mass customization."¹⁹

As Web-based digital maps progressively become the main portal through which to view the world and its places, the idea of applying mass personalization to maps comes within reach. It is now conceivable to develop mapping platforms that generate personalized maps not only for a specific task but for a specific individual, taking into account the individual's experience, behavior, knowledge, and particular viewpoint. Surprisingly, while many online products and services have been personalized over the past decade, dig-

ital maps are still fundamentally untouched by mass customization.

Map Personalization So Far

Maps are complex cultural and technical objects that assemble multiple data sources, assumptions about the user, cartographic traditions and practices, and design choices. All elements that form a digital map can in principle be personalized to increase the usability, efficiency, and clarity of the map with respect to a task.²⁰ To personalize maps, useful information can be provided by the user through explicit or implicit feedback.¹⁰

"Explicit feedback" is the conscious selection of preferences in the interface, and any action that is explicitly aimed at expressing a preference about any element of the map, as in, say, changing the language and the default settings of the interface. By contrast, implicit feedback includes any data about the user who expresses a preference indirectly; for example, moving the mouse cursor and clicking on a geographic area expresses a form of interest in that region, while hiding a layer at the beginning of each session indicates lack of relevance. The user's location represents another instance of implicit feedback about what areas of the map are of particular interest. Implicit feedback can in principle be extracted from any data generated by the user, including activity on social media, instant messaging, purchases from online stores, and email messages.

The most important indicator that must be determined is the task the user is currently performing on the map. Common tasks performed on popular online services (such as Yahoo! Maps and Google Maps) include information retrieval, general exploration of a region of interest, and routing. Different maps suit different tasks, with respect to features, layers, and controls. While in some cases the user's intentions are easy to detect (such as typing a place name is likely to indicate an information retrieval and/or a routing task), many behaviors do not imply specific tasks and present a considerable interpretive challenge. Automatic task detection can be performed with many indicators, including user demographics, interaction logs, search history, and the user's context and its spatial, temporal,

social, and computational aspects.

Over the past 15 years, academic researchers, including us, have investigated ideas and techniques through which map personalization can be achieved, working in two complementary strands. On the one hand, the automated adaptation of the map is pursued to increase clarity and efficiency and reduce information overload by removing or highlighting features based on user preferences and current task. On the other, the area of recommendation has generated techniques to personalize search results and recommendations of hotels, restaurants, and other points of interest based on individual and/or collective preferences. These two strands overlap in that similar techniques can be used to adapt the map and personalize search results and recommendations.

In a pioneering work in 2000, Oppermann and Specht¹⁷ developed *Hippie*, a tool that presents museum information based on the context of use. It relies on a user model that represents the user's knowledge and interests, a domain model of the information being displayed, and a space model in which the interaction occurs. Brunato and Battiti⁴ devised the Personal Item Locator and General Recommendation Index Manager (PILGRIM), a recommender system that takes the user's location into account to rank Web pages. Although recommender systems have been widely adopted since the mid-2000s, little work has been done to increase their spatial awareness.²¹

Our work focuses on the use of implicit feedback to adapt the map content itself. The core assumption is that implicit feedback indicators (such as mouse movements and navigational behavior) can be used to infer user interests.¹³ A recurring cognitive issue, particularly in the context of mobile computing, is that of spatial information overload, or display of excessive amounts of information on the map, hindering, rather than helping, the user. CoMPASS is a GIS application that monitors user interaction to recommend groups of features (such as layers) to users.^{12,22} The MAPPER system generates maps containing specific features, taking into account the user's preferences and the computational context by implicitly monitoring the interactions of users when brows-


Figure 1. Alternative cartographic representations of University College Dublin. Maps a–c are from commercial services, and d–f are based on OpenStreetMap open data using different themes; visualization generated with GeoFabrik tools.



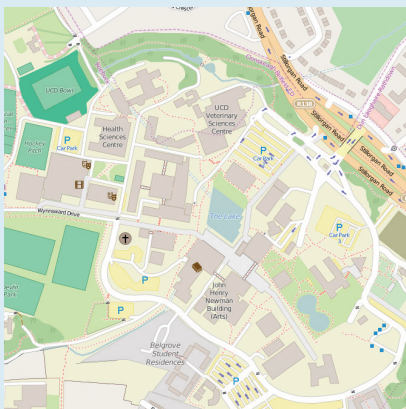
(a) Bing Satellite



(b) Google Maps



(c) Bing Maps



(d) OpenStreetMap



(e) OpenStreetMap



(f) OpenStreetMap

ing maps and inferring individual and group preferences.²³ This approach has been evaluated on a variety of map-based tasks,²⁴ showing increased efficiency in task completion; a similar approach was applied to detect the current task and adapt the map to it.¹⁴

In the RecoMap prototype,² we explored the possibility of computing interest scores for geographic features based on two complementary aspects—interaction (amount of interaction with a feature) and proximity (physical proximity to the feature) to generate personalized recommendations. A memory model simulates the decay of interest over time, assuming if the user does not interact with an object, the user’s interest in it is declining. Moreover, we investigated the possibility of integrating crowdsourced spatial data into the personalization analysis and utilizing Linked Open Data, a net-

work of inter-connected datasets, to increase the semantic structure of the geographic features.¹ Although this body of research initiated the theoretical and practical development of map personalization, new concepts, paradigms, and techniques await further investigation and evaluation.

This line of research promises to generate a plethora of commercial applications, greatly enriching current Web-mapping platforms. Since the mid-2000s, following increased bandwidth and more sophisticated Web browsers, a growing non-specialist mass market for Web maps has emerged, first on desktop computers, and more recently on GPS-enabled smartphones.^a In order to

review the state of the art of map personalization in products on the consumer market, it is useful to distinguish between manual and automated personalization, reflecting implicit and explicit feedback. Manual personalization allows users to modify aspects of the map, using preferences, bookmarking, and map editors. By contrast, automated personalization relies on implicit feedback to modify the map without user intervention, using data mining to model the user’s tastes and intentions; the table here lists personalization capabilities of popular, global, currently active mapping services.

Google Maps is the only service today that provides some automated personalization, tailoring the search results and ads based on the user’s search history and ratings, claiming to generate “a map for every person

a http://www.comscore.com/Insights/Blog/Map_Searches_Shift_from_Desktops_to_Smartphones

Personalization in Web maps, October 2015.

Product	Manual personalization	Automated personalization
Google Maps	Vector base map. In "My Places," users can set their home and work address. MapsEngine cloud tool can be used to create new maps by adding layers on a set of base maps.	Search results, recommendations, and advertisements are based on previous ratings and searches.
ArcGIS Online (ESRI)	Complex new maps can be created, combining base maps with user-defined layers. Advanced user-defined analytics available.	None.
Apple Maps	Vector base map. Users can bookmark locations in iCloud.	None.
OpenStreetMap (OSM Foundation)	OSM open data can be used to generate customized maps with a variety of dedicated tools (such as MapBox).	None.
HERE (Nokia)	Users can create "collections" of locations.	None.
MapQuest (AOL)	Users can save favorite locations and vehicles to improve routing. With "My Maps," they can save collections of locations and routes.	None.
Yahoo! Maps	Based on Nokia's HERE.	None.
ViaMichelin	In "My Michelin," users can bookmark locations, restaurants, and itineraries.	None.
OS Map Finder (U.K. Ordnance Survey)	Users can draw paths.	None.
Bing Maps (Microsoft)	Users can bookmark locations in "My Places."	None.

and place.”^b Other popular Web mapping products (such as ArcGIS Online and Yahoo! Maps) offer some manual personalization, typically in the form of bookmarks or editors to create and share new maps with user-provided data and visual styles. None of these products attempts to perform automated personalization.

Computational Challenges

The considerable increase in variety

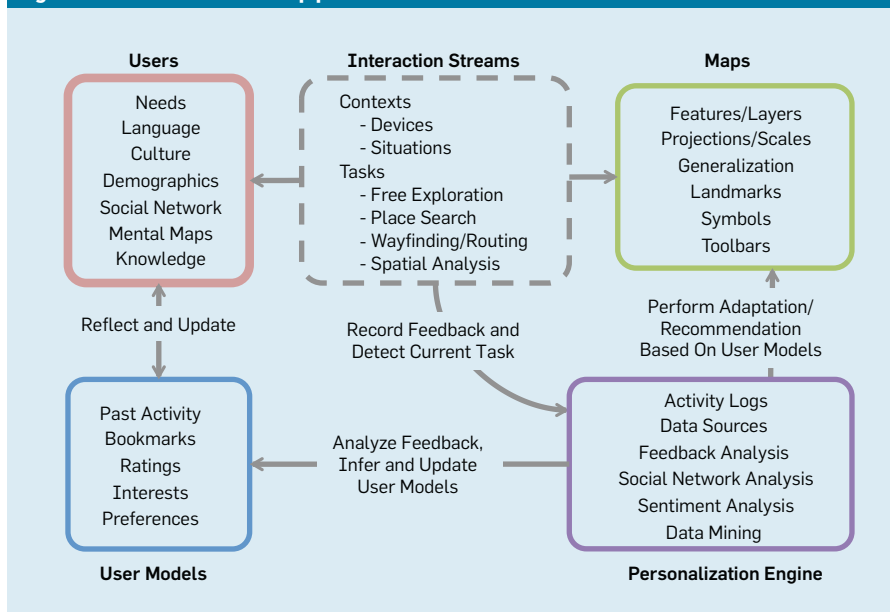
^b <http://google-latlong.blogspot.it/2013/05/meet-new-google-maps-map-for-every.html>

and volume of data provides a largely unexplored yet fertile basis from which to rethink maps. As we have shown, the field of map personalization is still at an early stage, with little research conducted or applied to commercial products. To further map personalization, several components must be integrated into a coherent conceptual framework (see Figure 2). A personalization engine must be able to perform multivariate feedback analysis on the many channels through which users express their spatial interests and preferenc-

es. By monitoring how users perform their tasks, the engine should be able to mine and extract meaningful patterns, inferring effective user models. Based on these models, the current task can be detected and trigger personalization in its two dimensions—adaptation and recommendation—at appropriate moments, unobtrusively supporting the user.

Many possibilities lie ahead to deliver more relevant, effective, and useful maps to Internet users. Given the same geographic data, such a system will tailor different maps for, say, a Japanese tourist in San Francisco and for an Italian ex-pat who lives and works in San Francisco. The tourist explicitly indicated an interest in architecture and a dislike for fast food. Her interaction shows implicitly an interest in historical areas, reflected in the map by increasing the prominence of history-themed museums. The map also captures that she was in the city before, displaying previously visited points of interest, facilitating spatial comprehension and way-finding, while at the same time emphasizing unknown areas of the city that feature notable buildings. Fast food restaurants remain hidden, unless she searches for them explicitly or when they might provide useful navigation landmarks. Before meal times, the map emphasizes restaurants that were recommended by her friends

Figure 2. A framework for map personalization.



who visited the city, while at night it increases the visibility of movie theaters, taking into account her interest in cinema.

By contrast, the Italian ex-pat specified an interest in music. As he navigates the city, the map tends to hide familiar tourist attractions while highlighting content neighboring his home-to-work commutes. Based on his interaction history, the system also captures an interest in the Mission District, which becomes more prominent and detailed on the map. As he uses a car to drive around the city, the public transport infrastructure fades into the background of the map. However, if his movement patterns match an efficient bus route, the system discretely suggests an alternative transportation option. Occasionally, the map emphasizes an unfamiliar neighborhood that presents a high density of music venues, inviting him to go beyond the borders of his daily routine. To make this vision real, several computational challenges lie ahead.

Real-time task detection and prediction. As maps are used in a variety of situations for different tasks, the system must be able to detect and predict them effectively. For this purpose, specific machine-learning techniques must be developed and optimized. In the collection of implicit feedback, relevant features include the user's spatial and temporal context, as well as search and interaction streams.

Spatial user modeling. Because of maps' spatial nature, personalization requires deep understanding of user behavior in space and time. Hence, the aggregation and interpretation of large numbers of noisy spatiotemporal trajectories containing GPS fixes, clicks, and search logs are essential for developing models of user behavior able to capture and predict recurring patterns and anomalies (such as sightseeing, as opposed to daily commuting). Recording, storing, and mining a large volume of spatiotemporal trajectories constitute an open research challenge.²⁵ Trajectories can traverse the geographic space, as well as other spaces, including mouse trajectories in a user interface. Spatial social network analysis can also illuminate the deep structures that influence interaction with maps.

Geographically weighted personalization. As maps assist users in the exploration and navigation of the geographic space, personalization should be able to tap the spatial variation in human activities. Development of geographically weighted techniques relies on the assumption that individual information needs and content relevance change both spatially and temporally. Such spatialization of users and content and their relationships increases the computational complexity of traditional personalization models. In turn, spatiotemporal user models would enable finer and more sensitive personalization of contents.

Geosemantic interoperability and data fusion. Map personalization requires aggregation and fusion of a range of heterogeneous sources of geographic information characterized by intrinsic uncertainty, vagueness, and rapid obsolescence, ranging from traditional government agencies to crowdsourced, volunteered datasets. In this sense, research in the context of the semantic Web and linked open data provides suitable representational tools to organize, store, explore, and retrieve spatiotemporal objects and data streams.⁹ Reliable mechanisms are needed to reference entities and perform identity resolution, reducing the friction caused by interoperability issues.³

Geoparsing and sentiment analysis. As a vast amount of geographic knowledge is expressed in natural language, natural language processing (NLP) is crucial for map personalization, extracting value from unstructured data. Geoparsing, or extraction of geographic information from natural language, is an open problem in NLP, intimately connected to word-sense disambiguation. Detection of affect, sentiment, and emotion in text is an emerging yet important aspect of interpreting user behavior, improving extraction and modeling of users' opinions about places.

Cognitive map design. Concepts and principles from cognitive map design¹⁵ can be applied to generate and validate alternative cartographic representations, providing an exciting opportunity to test cognitive theories against real scenarios on large numbers of users. From a technologi-

cal viewpoint, interactive maps have moved from a tile-based approach, in which maps are served through cached pre-rendered images, to a more flexible vector-based approach, in which the rendering occurs in real time in the client, providing the ideal platform for experimenting with alternative rendering choices and styles. Knowledge from spatial cognition could also be useful for producing better personalized maps and gain further insight into the human perception and understanding of the geographic environment.⁸

To achieve significant advances in these areas, academic researchers and commercial developers must tap the informational wealth produced by millions of users worldwide in their daily interactions mediated by online platforms, in which space and place are deeply intertwined with social, cultural, and economic processes. Due to the intrinsic complexity of these processes, map personalization needs help from thriving research areas in the context of big-data analytics. From a complementary perspective, NLP and sentiment analysis can be used to mine user-generated opinions about places to increase or decrease the emphasis of specific features. Advanced data-mining techniques are needed to extract meaning from noisy interaction logs. Beyond these computational steps, the challenges of map personalization are intrinsically multidisciplinary, harnessing ideas and tools from geographic information science, cartography, cognitive psychology, human-computer interaction, and software design.

Consequences

From a societal viewpoint, automated map personalization at a mass scale could have serious implications that should be responsibly taken into account, particularly by commercial developers whose products reach millions of users. Beyond the obvious concern for privacy, fostered by any surveillance-based technology, specific problems include the potential loss of a common representation of geographic realities. Personalized maps might result in what Internet activist Eli Pariser calls a "filter bubble," increasing social and cultural seg-

regation between groups of users.¹⁸ Likewise, personalized landmarks can be useful for increasing the clarity of maps but might also reduce the common semantic ground shared by the inhabitants of a geographical area. In this regard, Google Maps, currently the only commercial product to include some form of map personalization, presents several unresolved questions. The most conspicuous is the product's lack of transparency, making it difficult for users to understand why certain features are recommended over others. The user models generated by Google are black boxes not accessible to the users they are supposed to represent, and, more important, there is no visible "off" button to disable the personalization; even when logged out, the search results are still personalized in unclear ways based on cookies and the IP location of the user's machine.

A serious challenge for academic research in map personalization is the lack of realistic interaction datasets for evaluating novel systems and approaches. As private corporations are understandably reticent to share their map-interaction logs, the studies discussed here have limited evaluations, failing to reflect the complexity, noise, and variety of situations in real mapping applications, limiting their observation to small, artificial, and controlled contexts. A few large corporations (such as Google and Microsoft) attract the vast majority of online map users and are thus in a privileged position to unobtrusively devise and evaluate proprietary techniques on large groups of users on a variety of tasks, interpreting their behavior as implicit feedback. For this reason, academic research must either focus on well-defined cognitive, computational, and cartographic aspects of map personalization that can be convincingly evaluated or work in close partnership with map providers. No progress in map personalization can be assessed in the absence of rigorous measures to quantify the effectiveness of techniques, algorithms, and models.

Conclusion

Development of personalized maps has important applications in many domains. To date, the focus has been

on commercial applications; for example, Google has been exploring location-based advertising, trying to maximize the relevance and profitability of ads that have a strong spatial component. Similarly, most research focuses on efficiency, reducing information overload, increasing clarity, and helping users complete tasks more quickly or with lower cognitive load, as in, say, decision making, information retrieval, and routing. However, exciting possibilities also exist beyond increased efficiency. Personalized maps need not reinforce users' biases and limited perspectives but can be designed to operate in the opposite way, attracting attention to the unknown and unfamiliar and promoting diversity, serendipity, and discovery. In education, self-adaptive maps could support students, tailoring maps to different learning styles and backgrounds. It is reasonable to expect map personalization could trigger a quiet but deep reconfiguration of familiar maps, leading to unexpected changes in the way we perceive and imagine the world around us.

Acknowledgments

The research presented here was funded by a Strategic Research Cluster grant (07/SRC/I1168) by Science Foundation Ireland under the National Development Plan. C

References

1. Ballatore, A. and Bertolotto, M. 2011. Semantically enriching VGI in support of implicit feedback analysis. *Web and Wireless Geographical Information Systems*, K. Tanaka, P. Fröhlich, and K.-S. Kim, Eds. LNCS, Vol. 6574 (Kyoto, Japan, Mar. 3–4). Springer, Berlin, Germany, 78–93.
2. Ballatore, A., McArdle, G., Kelly, C., and Bertolotto, M. RecoMap: An interactive and adaptive map-based recommender. In *Proceedings of the 25th ACM Symposium on Applied Computing* (Sierre, Switzerland, Mar. 22–26). ACM Press, New York, 2010, 887–891.
3. Ballatore, A., Wilson, D., and Bertolotto, M. A survey of volunteered open geo-knowledge bases in the semantic Web. In *Quality Issues in the Management of Web Information*, G. Pasi, G. Bordogna, and K. Jain, Eds. Intelligent Systems Reference Library, vol. 50. Springer, Berlin, Germany, 2013, 93–120.
4. Brunato, M. and Battisti, R. PILGRIM: A location broker and mobility-aware recommendation system. In *Proceedings of the IEEE International Conference on Pervasive Computing and Communications* (Fort Worth, TX, Mar. 23–26). IEEE, Piscataway, NJ, 2003, 265–272.
5. Gartner, G., Bennett, D., and Morita, T. Towards ubiquitous cartography. *Cartography and Geographic Information Science* 34, 4 (2007), 247–257.
6. Haklay, M., Ed. *Interacting with Geospatial Technologies*. John Wiley & Sons, Chichester, U.K., 2010.
7. Haklay, M., Singleton, A., and Parker, C. Web mapping 2.0: The neogeography of the GeoWeb. *Geography Compass* 2, 6 (Nov. 2008), 2011–2039.
8. Hirtle, S.C. Geographical design: Spatial cognition and geographical information science. *Synthesis Lectures on Human-Centered Informatics* 4, 1 (2011), 1–67.
9. Janowicz, K., Scheider, S., Pehle, T., and Hart, G. Geospatial semantics and linked spatiotemporal data: Past, present, and future. *Semantic Web* 3 (2012), 321–332.
10. Kelly, D. and Teevan, J. Implicit feedback for inferring user preference: A bibliography. *ACM SIGIR Forum* 37, 2 (Fall 2003), 18–28.
11. Lynch, K. *The Image of the City*. MIT Press, Cambridge, MA, 1960.
12. Mac Aoidh, E. and Bertolotto, M. Improving spatial data usability by capturing user interactions. In *The European Information Society: Leading the Way with Geo-information*, S.I. Fabrikant and M. Wachowicz, Eds. Springer, Berlin, Germany, 2007, 389–403.
13. Mac Aoidh, E., Bertolotto, M., and Wilson, D.C. Analysis of implicit interest indicators for spatial data. In *Proceedings of the 15th Annual ACM International Symposium on Advances in Geographic Information Systems* (Seattle, WA, Nov. 7–9). ACM Press, New York, 2007, 1–4.
14. Mac Aoidh, E., McArdle, G., Petit, M., Ray, C., Bertolotto, M., Claramunt, C., and Wilson, D. Personalization in adaptive and interactive GIS. *Annals of GIS* 15, 1 (2009), 23–33.
15. Montello, D. Cognitive map-design research in the twentieth century: Theoretical and empirical approaches. *Cartography and Geographic Information Science* 29, 3 (2002), 283–304.
16. Negroponte, N. *Being Digital*. Random House, New York, 1995.
17. Oppermann, R. and Specht, M. A context-sensitive nomadic exhibition guide. In *Handheld and Ubiquitous Computing*, P. Thomas and H.-W. Gellersen, Eds. LNCS, Vol. 1927. Springer, Berlin, Germany, 2000, 127–142.
18. Pariser, E. *The Filter Bubble: How the New Personalized Web Is Changing What We Read and How We Think*. Penguin, London, U.K., 2012.
19. Piller, F.T. and Tseng, M.M., Eds. *Handbook of Research in Mass Customization and Personalization, Volume 1*. World Scientific, Singapore, 2010.
20. Reichenbacher, T. Adaptive concepts for a mobile cartography. *Journal of Geographical Sciences* 11, 1 (2001), 43–53.
21. Stiller, C., Ros, F., and Ament, C. Towards spatial awareness in recommender systems. In *Proceedings of the International Conference for Internet Technology and Secured Transactions* (London, U.K., Nov. 9–12). IEEE, Piscataway, NJ, 2009, 1–7.
22. Weakliam, J., Bertolotto, M., and Wilson, D. Implicit interaction profiling for recommending spatial content. In *Proceedings of the 13th Annual ACM International Workshop on Geographic Information Systems* (Bremen, Germany, Nov. 4–5). ACM Press, New York, 2005, 285–294.
23. Weakliam, J., Wilson, D., and Bertolotto, M. Personalising map feature content for mobile map users. In *Map-based Mobile Services: Theories, Methods and Implementations*, L. Meng, A. Zipf, and T. Reichenbacher, Eds. Springer, Berlin, Germany, 2008, 125–145.
24. Wilson, D., Bertolotto, M., and Weakliam, J. Personalizing map content to improve task completion efficiency. *International Journal of Geographical Information Science* 24, 5 (2010), 741–760.
25. Zheng, Y. and Zhou, X., Eds. *Computing with spatial trajectories*. *Database Management & Information Retrieval*. Springer, Berlin, 2011.

Andrea Ballatore (aballatore@spatial.ucsb.edu) is a postdoctoral researcher and research coordinator in the Center for Spatial Studies at the University of California, Santa Barbara.

Michela Bertolotto (michela.bertolotto@ucd.ie) is a senior lecturer in the School of Computer Science and Informatics at the University College Dublin, Ireland.

© 2015 ACM 00010782/15/12 \$15.00



Watch the authors discuss their work in this exclusive *Communications* video. <http://cacm.acm.org/videos/personalizing-maps>

Copyright of Communications of the ACM is the property of Association for Computing Machinery and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.