Semantic Signatures for Places Of Interest

Krzysztof Janowicz
STKO Lab
Department of Geography
University of California, Santa Barbara, US
janowicz@ucsb.edu

Benjamin Adams
Centre for eResearch
Department of Computer Science
University of Auckland, NZ
b.adams@auckland.ac.nz

Search was always spatial. Searching and ranking text documents, for instance, is typically based on vector space models where similarity is calculated as the cosine of the angle between the term vectors. According to information foraging theory, such documents, say Web sites and their pages, contain information patches and their spatio-temporal properties are exploited by informavores to select a promising path based on the information scent [1]. In a broader sense, even typing in a search query, e.g., a sequence of keywords, is spatial. Terms closer to another are more likely to form meaningful n-grams. Of course, there is also spatial variation in what is being searched. A user living in a region with good public transit is more likely to search for bus routes and time tables [2]. Changing perspective, the dissemination and diffusion of information is also known to follow spatio-temporal patterns [3]. Finally, geo-fencing takes the act of querying out of the loop by *pushing* notifications to a device that enters or leaves a store, event, or another digitally bounded area. Examples range from sending coupons to users that walk by a store or alerting users about potential theft if their car leaves the parking area without them.

In most cases, however, when we refer to *spatial search* we mean the fact that the user's geographic location provides important contextual cues to improve the relevance ranking between the query and the objects under consideration, e.g., Places Of Interest (POI). *Spatial contextual awareness* is just one of many contextual cues; other examples include the user's profile, navigation history, device type, and so forth. However, location is widely considered to be highly indicative of the user's intent [4]. Simplifying, a search engine will return a nearby coffee shop when queried for *coffee* instead of a more distant one or an Web page about the history or politics of coffee. What was true once only for the search on mobile devices is now also common practice for the desktop. Additionally, some systems support simple constraints on pre-defined attributes, e.g., place type, wheelchair access, wifi availability, or excluding localities that are currently closed based on the time the query is posed. A new category of applications centered around the idea of intelligent personal (digital) assistants relates events to locations, e.g., by showing traffic data while the user is on the way to work. Summing up, *location* matters for *search,* is typically represented by a device's geofix, and by *spatial* we most often mean *nearby*.

This, however, does not necessarily mirror the rich human experience of geographic space and particularity the interaction with *places*. One could now argue that the focus should be on richer geometric representations of these places by polylines and polygons instead of simple centroids. This would also enable topological queries such as for rivers that flow through a parks. While this is certainly a valid point, it again reduces the richer notion of place to spatial footprints and search to geometrical relations, pre-defined attributes, and top-down defined place types, e.g., Jazz club. **Instead, we need softer queries on smarter data and more expressive, non-reductionistic computational models of place**. For instance, to start with a simple scenario, one can query for hotels near a given landmark or maybe even away from an airport, but it is not possible to query for a hotel in a central but quiet location. This is for two reasons: (i) negative queries are difficult to evaluate and rank; clearly the user does not want to stay somewhere in nowhere, so how far away from noise should the hotel be; (ii) noise is a *neighborhood-level* attribute and cannot be easily reduced by excluding hotels that are nearby a pre-defined set of place types, e.g., schools and concert halls. Today, such queries are handled by

returning those hotels for which existing reviews complain about noise and then letting the user manually picky from the remaining hotels. A better approach would to understand neighborhoods by the distribution of POI, traffic flow patterns, latent characteristics, and so forth. Simplifying, one could exclude hotels that are nearby places of types that are known to generate noise. Instead of having to rely on a pre-defined set of POI types as specified by schema.org, Fourquare, Yelps, and so forth, i.e., by extensionally enumerating them as {nightclub, disco, bar, airport, concert hall, school, …}, one could generate the *noisy places* type bottom-up by identifying types of places that show peak (or continued) activity patterns during the evenings. To intuitively understand why such an approach is more robust, consider the example of nightclubs, bars, and schools. Clearly, one would assume that a hotels in the direct vicinity of either bars, nightclubs, or schools, will be exposed to noise. However, typically a guest staying during the workweek can safely ignore the nightclubs as a factor, while bars are noisy almost every evening throughout the week. Schools, in contrast, have their activity peaks during typical working hours and, thus, would less likely disturb a business traveler.

In other words, place types can be defined and combined bottom-up based on the temporal behavior of humans towards them. This perspective opens up new lines of research such as estimating place types for unlabeled places based on check-in patterns as well as improving state-of-the-art geolocating services that match a user's spatial location to potential places. This is an important task, as *platial* information is semantically richer than just spatial proximity. For instance, standing in front of (or nearby) a food truck does not entail that the user will buy something or is even waiting in line. In contrast, checking-in is a active commitment to being at a place. Time, of course, is not the only characteristics that can be exploited this way. As analogy to spectral signatures in remote sensing, we have introduced *semantic signatures* that are made out of temporal, spatial, and thematic *bands* [5,6] and enable us to classify places and regions based on their unique characteristics. In a nutshell, certain types of places can be distinguished based on a single band (e.g., check-in peaks during the day) alone, while others require multiple bands that jointly form the unique signature of this place type.

To return to the previously introduces example of searching for hotels in a central but quiet location, one could argue that the temporal bands alone cannot sufficiently address the noise problem. After all fire stations, police stations, or hospitals will have regular temporal patterns throughout the day. Thematic bands derived via topic modeling from user reviews and place descriptions, however, will reveal this information by mentioning emergency sirens, noise, and so forth. In fact, even non-georeferences everyday language is geo-indicative to a degree where it allows us to characterize regions or neighborhoods or infer were certain phrases are more likely to be uttered [7]. Similar to the temporal case, such thematic bands reveal interesting latent topics and provide bottom-up categories. For instance, people do not just got to any bar, they favor certain types of bars that cannot be distinguished just by properties such as wifi access or price level alone. For instance, thematic bands can reveal that certain people go to bars that offer *pub quizzes*. Even the > 400 POI categories in foursquare are not granular enough to make such distinctions (and classical search would rely on the appearance of the exact pub quiz label). Finally, one could ask about the spatial extent of a quiet neighborhood (in the sense defined above). Spatial bands (and signatures derived from them) allow us to characterize places based on the spatial relation to other places of the same or different types, thereby exploiting the fact that places such as bars co-occur with other bars, nightclubs, and restaurants, while police stations, post offices, and so forth are regularly distributed and unlikely to co-occur.

Finally, utilizing such signatures for a semantics-enabled spatial search begs the question of how to weight the individual components, i.e., assuming a linear combination of space, time, and topic, is it more important to match a spatial extent or a (semantically expanded) theme? To finish with our starting example, is *central* more important than *quiet*?

References

[1] Pirolli, P., & Card, S. (1999). Information foraging. Psychological review, 106(4), 643-675.

[2] Backstrom, L., Kleinberg, J., Kumar, R., & Novak, J. (2008, April). Spatial variation in search engine queries. In Proceedings of the 17th international conference on World Wide Web, pp. 357-366. ACM.

[3] T. Hägerstrand (1967) Innovation Diffusion as a Spatial Process. University of Chicago Press .

[4] Teevan, J., Karlson, A., Amini, S., Brush, A. J., & Krumm, J. (2011). Understanding the importance of location, time, and people in mobile local search behavior. In Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services (pp. 77-80). ACM.

[5] Ye, M., Janowicz, K., Mülligann, C., & Lee, W. C. (2011). What you are is when you are: the temporal dimension of feature types in location-based social networks. In Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (pp. 102-111). ACM.

[6] McKenzie, G., Janowicz, K., Gao, S., Yang, J. A., & Hu, Y. (2014). POI Pulse: A Multi-Granular, Semantic Signatures-Based Approach for the Interactive Visualization of Big Geosocial Data. Cartographica: The International Journal for Geographic Information and Geovisualization, The University of Toronto Press.

[7] Adams, B., & Janowicz, K. (2012). On the Geo-Indicativeness of Non-Georeferenced Text. The 6th International AAAI Conference On Weblogs And Social Media (pp. 375-378). AAAI.