Computers, Environment and Urban Systems 54 (2015) 1–13

Contents lists available at ScienceDirect

Computers, Environment and Urban Systems

journal homepage: www.elsevier.com/locate/ceus

Where is also about time: A location-distortion model to improve reverse geocoding using behavior-driven temporal semantic signatures

Grant McKenzie *, Krzysztof Janowicz

The STKO Lab, Department of Geography, University of California, Santa Barbara, CA, USA

ARTICLE INFO

Article history:
Received 19 November 2014
Received in revised form 21 May 2015
Accepted 22 May 2015
Available online 12 June 2015

Keywords:
Reverse geocoding
Point of interest
Semantic signature
Geosocial check-in
Time

ABSTRACT

While geocoding returns coordinates for a full or partial address, the converse process of reverse geocoding maps coordinates to a set of candidate place identifiers such as addresses or toponyms. For example, numerous Web APIs map geographic point coordinates, e.g., from a user's smartphone, to an ordered set of nearby Places Of Interest (POI). Typically, these services return the k nearest POI within a certain radius and measure distance to order the results. Reverse geocoding is a crucial task for many applications and research questions as it translates between spatial and platial views on geographic location. What makes this process difficult is the uncertainty of the queried location and of the point features used to represent places. Even if both could be determined with a high level of accuracy, it would still be unclear how to map a smartphone's GPS fix to one of many possible places in a multi-story building or a shopping mall. In this work, we break up the dependency on space alone by introducing time as a second variable for reverse geocoding. We mine the geosocial behavior of users of online location-based social networks to extract temporal semantic signatures. In analogy to the notion of scale distortion in cartography, we present a model that uses these signatures to distort the location of POI relative to the query location and time, thereby reordering the set of potentially matching places. We demonstrate the strengths of our method by evaluating it against a purely spatial baseline by determining the Mean Reciprocal Rank and the normalized Discounted Cumulative Gain. Our method performs substantially better than said baseline.

1. Introduction and motivation

Translating back and forth between spatial and placial representations of location is a crucial task underlying many research questions, applications, and systems. Geocoding, for instance, is the process of assigning corresponding geographic coordinates to other types of structured geographic identifiers such as addresses. The converse process, called reverse geocoding, assigns place identifiers, such as toponyms, to geographic coordinates. More specifically, it maps a geometry in the sense of OGC's Simple Feature model to an ordered set of candidate place identifiers. Typically, the Euclidean distance between the query coordinates and the point-feature representation of the candidate places is used to establish a relevance ranking. To successfully match a user's location to a visited place, new geosocial approaches also consider popularity, e.g., how many users checked-in or wrote reviews about a place. Additionally, many (reverse) geocoding systems consider place hierarchies and granularity.

The following queries nicely illustrate the difference between a spatial and placial perspective as well as the arbitrariness of relying on point coordinates for the query and the candidate places alone. While not a reverse geocoder in the strict sense, the Flickr *flickr.places.findByLatLon* API call (Flickr, 2014) returns place IDs given a lat/lng coordinate and accuracy value. This allows users to find photos for particular places. The API *rounds up* to the nearest place type, i.e., it returns a city ID for street-level coordinates rather than returning a street or building. Latitudes and longitudes are truncate to three decimal points. In each case shown below, the query coordinates represent the same fix at the Griffith Observatory in Los Angeles. However, the query is run with different accuracy levels where 16 corresponds to the street level, 11 to the city level,

* Corresponding author.
E-mail addresses: grant.mckenzie@geog.ucsb.edu (G. McKenzie), janowicz@ucsb.edu (K. Janowicz).

http://dx.doi.org/10.1016/j.compenvurbsys.2015.05.003
0198-9715/

and 7 to the county level. The respective responses from the Flickr API are as follows.[1]


```
<placeslatitude = "34.118341"longitude
 = " − 118.300458" accuracy = "16" total = "1" >
 <placeplace_id = "HqDLYDJTUb8XihsYDg" woeid
 = "23511984" latitude = "34.125" longitude
 = " − 118.306" […] place_type
 = "neighbourhood" place_type_id = "22" timezone
 = "America/Los_Angeles" name=
   Hollywood United, LosAngeles, CA,
   US, UnitedStates" woe_name
 = "HollywoodUnited"/ >
 </places >
<placeslatitude = "34.118341" longitude
 = " − 118.300458" accuracy = "11" total = "1" >
 […] latitude = "34.146" longitude
 = " − 118.248" […] place_type
 = "locality" place_type_id = "7" name
 = "Glendale,California,UnitedStates" […]" / > […]
<placeslatitude = "34.118341" longitude
 = " −118.300458" accuracy = "6" total = "1" >
 […] place_type = "county" place_type_id
 = "9" […] name = "Los Angeles County,California,
 UnitedStates" […] / > […]
```


The fact that even small differences in spatial accuracy may have strong impacts, e.g., on routing choices, has been demonstrated in the literature before (Bowling & Shortridge, 2010). What makes the example above interesting is the place hierarchy. Hollywood is a district of Los Angeles, while Glendale is a city in Los Angeles County. From a human-centered *placial* perspective, one would assume the queries to return Hollywood (in fact, it should be the Los Feliz neighborhood), Los Angeles, and finally Los Angeles County. Instead the neighboring city of Glendale is returned for the city-level accuracy query, thereby breaking the expected hierarchical composition of places. From a computation-centric *spatial* perspective Glendale is returned by the Flickr API simply because its centroid representation it closer to the query location than the centroid of Los Angeles.

The arbitrariness and imprecision of point-feature representations as well as the effect of missing topological relations also strikes on the level of small-scale features such as Places Of Interest (POI).[2] Fig. 1 illustrates a common issue. First, the resort marker (A) is placed at the entrance to the parking lot. While this may be acceptable, other POI databases place it at the center of the building which is nearly 150 m away. Second, the lounge is *inside* the resort but its marker (B) is shown over 100 m away from the resorts marker. As most reverse geocoders rely on distance alone, such differences will lead to substantially different and often misleading results, e.g., when suggesting a user's check-in location.

As the omnipresence of location-enabled mobile devices increases, more robust, accurate, context-aware, and data-rich geolocation services are required. Today, the ability to link spatial coordinates to an actual place has become essential in many aspects of our everyday lives including navigation applications, place recommendation, location-based advertising, and critical infrastructure. It is interesting to note that the challenge is not one of more accurate GNSS and Wi-Fi-based positioning systems (WPS) alone. The information that a person checked-in or is present at a place is *semantically richer* than the spatial data alone. To give a concrete example, the fact that a person is standing in front of a *food truck* is substantially different from the fact that a person checked-into the *food truck* and is likely to order something. Placial information is more than just spatial proximity.

Commercial companies such as *Google* as well as open source platforms like *GeoNames* have made names for themselves offering application programming interfaces (APIs) and web services that allow both developers and consumers to query gazetteers and POI databases using geographic coordinates as input. With the increase in user-generated geo-content, new services such as *Foursquare* and *Yelp* have emerged allowing anyone with a location-enabled mobile device to contribute or update the location of an entity in a crowd-sourced system. It is important to note that while these systems involve the contribution of geo-content from individual users, there is still some discussion as to whether or not they fit into the category of *Volunteered Geographic Information* (Harvey, 2014; McKenzie & Janowicz, 2014). Previous work on POI matching has shown that the median distance of a single POI between different geolocation service providers is 62.8 m apart and can reach up to several hundreds meters under extreme circumstances (e.g., for a golf course) (McKenzie, Janowicz, & Adams, 2014). Fig. 2 (left) illustrates this fact by showing the position of markers from five major services. While this offset may not be a substantial issue in rural areas due to their low POI density, it will cause substantial problems for geolocation services (e.g., check-in services) in high-density urban areas.

The task of determining the place an individual is visiting based on coordinates gathered from their mobile device becomes more difficult given the uncertainty associated with each POI in the dataset. That is, selecting the nearest POI to a user's location becomes an artifact of the arbitrary point-coordinate representation of nearby POI. Leaving the actual POI locations aside, another facet of uncertainty plagues traditional geolocation services, namely the positional accuracy of a location-enabled device. While most devices make use of a range of positioning technologies (e.g., GNSS, WPS, Cellular Network), each of these technologies has its own issues related to accuracy, imparting a level of uncertainty on any device location. Therein lies one of the problems facing traditional geolocation services such as reverse geocoders. Given the aforementioned sources of uncertainty, how can a geolocation service be expected to accurately predict a POI from geographic coordinates? An example of this challenge is shown in Fig. 2 (right). A number of POI are shown on the map along with their associated positional uncertainty. Additionally, the red pin shows the most probable location of a mobile device and it's two-dimensional depiction of uncertainty.

2. Research contribution and example scenario

Clearly, relying on geographic coordinates alone to infer a place based on a user's mobile device position is not sufficient. However, there are other contextual clues that can be taken into account. Time is one such clue and in contrast to many other contextual information it is readily available with every position fix. Current reverse geocoding services solely exploit geographic location while in reality human behavior dictates that approximately the same location in geographic space can serve a variety of purposes at different times of the day or days of the week. The motivation for visiting a specific city block on a Tuesday morning is considerably different than visiting that same block on a Saturday night. While the geographic coordinates determined by one's location-enabled mobile device may be temporally-agnostic, the *probability* of conducting an activity at a nearby place is not.

In fact, place categories are implicitly defined by time. For instance, the likelihood of being at the *Department of Motor Vehicles* on a Sunday

[1] The remainder of the paper will use data from the location-based social network Foursquare.

[2] Frequently also referred to as *Points Of Interest*.