Title

Spatial Search, Position Papers

Permalink

https://escholarship.org/uc/item/0h014085

Author

Center for Spatial Studies, UCSB

Publication Date

2014-12-01

Position Papers, 2014 Specialist Meeting—Spatial Search

TABLE OF CONTENTS

Predictive Spatial Search

MOHAMED ALI
Associate Professor, Institute of Technology
Center for Data Science
University of Washington, Tacoma
Microsoft Bing Maps
Email: mhali@uw.edu

with

ABDELTAWAB M. HENDAWI
Center for Data Science
Institute of Technology
University of Washington, Tacoma
Email: hendawi@uw.edu

MARTINE DE COCK
Center for Data Science
Institute of Technology
University of Washington, Tacoma
Email: mdecock@uw.edu

ANKUR TEREDESAI
Center for Data Science
Institute of Technology
University of Washington, Tacoma
Email: ankur@uw.edu

Abstract

In a typical spatial search problem, (mobile) users search for (stationary or mobile) entities that have spatial attributes. The user's current location and/or the entities' locations are considered to assess the relevance of the search result. On one hand, we believe that the user's future location is more relevant to the search result than the current location. Hence, we study spatial search queries under predictive models of user locations. On the other hand and with the ubiquity of hand held devices, most users do not utilize the full power of spatial search and they do not know what to search for. Hence, we introduce a framework to answer the question: "given a user's current and predicted locations, what would the user be interested in searching for and seeing as a query result?" More specifically, we propose a *predictive* spatial search approach that continuously monitors the user's current location to: (1) predict the user's future location and integrate this prediction efficiently in the spatial search query processing pipeline, and (2) predict the search keywords that are of relevance to the user, given the user's location and context. The second type of prediction leverages the knowledge of existing search engines about the behavior of a global set of spatial search users and social media users. Such two-phase prediction capability will enable search engines to "pre-search" on behalf of their users and, thereby, leading to gains in user experience, search accuracy, and communication costs.

Introduction

Commercial search engines recognize the value of tracking the users' location to provide a better service. The typical straightforward approach is to upload the user's location every time a search query is issued. However, the vision presented in [1] proposes a novel approach that utilizes the power of data streaming systems to track the users' locations based on a *route-selection preference policy*. In this approach, the system accepts (voluntarily) the user's declared route preferences along with other context-based information. This information helps the system "*predict*" the user's location even under the absence of the real-time location feed, survive offline periods, reduce communication cost and predict the user's future location with acceptable accuracy. The next set of challenges is to identify the issues that can help integrate spatial search capabilities with such predictive data streaming systems. This integration is the theme of our vision.

Consider the scenario that a user is driving on the highway and searches for a restaurant. It may be desirable and more relevant to direct the user to his favorite restaurant chain (or similar such chains that are still several minutes ahead from his current location) than to simply identify a nearby less favorite restaurant chain. It is even better to predict that this user is, for example, heading to a football game. Consequently, we direct the user to the parade, events and activities next to the stadium that the user is *not* aware of. In this scenario, we notice two main characteristics: first, the search engine is geared up to perform a spatial search around a non-deterministic (yet to be predicted) future location ahead of time. Second, the search engine aims at discovering and recommending personalized search topics on a per user basis. Ranking of the search result needs to be performed spatiotemporally taking into consideration the probability of the user's expected location along the future timeline.

Currently, users rely on multiple data sources such as live feeds from social media augmented with organic search to discover related events/facilities at the destination. An advantage of the proposed predictive spatial search frameworks is to help consolidate search results better across various faceted search channels. We believe that future geospatial search engine will inevitably adopt the "we know where you will be and *we search before you search*" principle. It is the proactiveness of search engines that will shape the future of spatial search. The search engine is the user's agent that continuously (1) adjusts the search result according to the user's predicted destination and (2) interacts/socializes with neighboring people/agents to highlight interesting topics that the user is totally unaware of.

Predictive Trees: An Index for Predictive Queries on Road Networks

Practical experience tells that it is absolutely a myth to assume that commercial search engines know everything about the user's past locations. Lots of research utilize manufactured databases of historical trajectories and apply machine learning techniques to predict the user's future location. These trajectory databases are usually collected by researchers or volunteers for research purposes. Yet, from a practical perspective, and due to privacy concerns [2], the user location is revealed on a session basis such that each session is no longer than few minutes. Commercial search engines care about users' privacy. Consequently, techniques that assume full knowledge of the user's behavior over extended periods of time are not considered practical in our approach. We propose a new index structure, the predictive tree [3, 4], that enables the evaluation of predictive queries [5] in the absence of the objects' historical trajectories. Based solely on the connectivity of the road network graph and assuming that the object follows the shortest route to destination, the predictive tree determines the reachable nodes of a moving object within a specified time window T in the future. Moreover, predictive trees utilize every additional piece of information and enhance the probability assignment of the predicted location as more trajectory data becomes available on the user.

The predictive tree: (1) provides a generic infrastructure for answering the common types of predictive queries including predictive point, range, KNN, and aggregate queries, (2) updates the probabilistic prediction of the object's future locations dynamically and incrementally as the object moves around on the road network, and (3) provides an extensible mechanism to customize the probability assignments of the object's expected future locations, with the help of user defined functions. In our ongoing effort, we leverage predictive trees to support spatial search and integrate this work with predictive data streaming systems.

References

[1] Mohamed Ali, Badrish Chandramouli, Balan Raman, and Ed Katibah. Spatio-Temporal Stream Processing in Microsoft StreamInsight. *IEEE Data Eng. Bull*. 33(2): 69–74 (2010).

[2] From GPS and Virtual Globes to Spatial Computing—2020: The Next Transformative Technology. A Community Whitepaper resulting from the 2012 CCC Spatial Computing 2020 Workshop.

[3] A. M. Hendawi, J. Bao, and M. F. Mokbel. Predictive Tree Source Code and Sample Data. URL:http://www-users.cs.umn.edu/~hendawi/PredictiveTree/, Aug. 2014.

[4] A. M. Hendawi and M. F. Mokbel. Panda: A Predictive Spatio-Temporal Query Processor. In ACM SIGSPATIAL GIS, 2012. [5] A. M. Hendawi and M. F. Mokbel. Predictive Spatio-Temporal Queries: A Comprehensive Survey and Future Directions. In MobiGIS, California, USA, Nov. 2012.

The Search for Places as Emergent Aggregates

ANDREA BALLATORE
Postdoctoral Researcher
Center for Spatial Studies
University of California, Santa Barbara
Email: aballatore@spatial.ucsb.edu

Searching for places is the most popular geographic online task, in which names and categories are used as the main referents to locate places in the geographic space. By typing "hotels in Santa Barbara" in any popular search engine, the user expects a list of places matching a category ("hotel") contained within another place called "Santa Barbara." Answers to such a query can be generated by relying on a gazetteer containing some form of spatial footprint for the symbol "Santa Barbara," a database of points-of-interest categorized as "hotels" and, indeed, some strategy to compute the relevance of the potential results. This approach satisfies well-defined information needs, but fails to account for more complex, nuanced, fuzzy, and yet cognitively intuitive questions about the town. What places are similar to Santa Barbara with respect to its general atmosphere—but perhaps less expensive? What other towns in Southern California offer a comparable array of amenities? What tourist areas in Italy provide a similar combination of mountain-related and marine activities? Our current computational models of place search do not seem to provide easy answers.

Our intimate familiarity with place clashes with the difficulty of dealing with it computationally. Because of its centrality in human cognition and culture, the notion of place is unsurprisingly characterized by high polysemy, strong context-dependence, and innumerable metaphorical uses, carving social meanings from neutral, unbounded spaces (Agnew, 2011). The intellectual prominence of place has waxed and waned over time, being obscured for centuries by more abstract notions of space, and making a reappearance in recent decades (Casey, 1997). Being intensely debated in the social sciences and the humanities, place and its representations have now become an active research frontier in geographic information science (Goodchild, 2011). In this area, a central concern is that place names are often ambiguous, vague, and vernacular. Place categories are culturally-dependent, arbitrary, and inconsistently applied. More strikingly, places are implicitly assumed to have a name and to fit, at least to some degree, known categories. Current efforts focus on more sophisticated place-name interpretation in text documents (Purves, and Jones, 2011), new reference theories tailored to place (Scheider and Janowicz, 2014), and semantically more expressive gazetteers (Keßler et al., 2009).

In a complementary approach, I advocate a view of *place as an aggregate* of objects and processes that interact at a given scale, inter-locked by spatial colocation. This view of place relies on the discovery of *implicit relations*, and not on some explicit labels assigned by an observer. The approach relies on some assumptions: Places are inescapably multi-faceted (comprising diverse processes), they are socially constructed (emerging as the result of human agency and practices), relational (emerging in a context, not in a vacuum), scale-dependent (different places exist at different scales), and they are dynamic (emerging, changing, and ultimately disappearing). Following Thrift (1999), I regard places as emergent entities in a complex, non-linear system, and they appear as assemblages of heterogeneous things that meet in space and time. Place can be fruitfully viewed through a holistic lens, emphasizing its contextuality and inherent interconnectedness, rather than as an object in isolation (Ballatore et al., 2012). In practical terms, this approach aims at supporting multi-faceted, context-dependent aggregate search, going beyond the current forms of name-based search for well-defined, individual places. In this sense, places can be searched for on the basis of their emergent distributional characteristics, rather than in an arbitrary, crisp categorization (e.g., *city* or *town*). For example, a Japanese tourist in San Francisco might search for places in which architectural landmarks co-occur with museums and galleries, or for places that, as an aggregate, present similar characteristics to Shibuya, the Tokyo shopping district she is familiar with.

	Text Information Retrieval	Place-as-aggregate Search
Vector space	Set of text documents	Geographic space
Vector	A document as a sequence of words	A place as an aggregate of spatially located objects
Dimensions	Words (high dimensionality)	Characteristics of objects (dimensionality defined by application, potentially very high)
Index	Sparse document-word matrix	Sparse object-to-object colocation matrix
Search	Weighted keyword matching, topic models, similarity	Colocation queries, query-by-place

Table 1: Place search overview

Computationally, this approach to place search can be modeled in analogy to traditional information retrieval in a vector space model, as summarized in Table 1. Given a geographic space, treated as a corpus containing a potentially infinite number of place-as-aggregate, the proposed approach has a number of opportunities and challenges to be tackled. The efficient computation of spatial colocation at a large scale—the identification of categories of objects that co-occur spatially (and temporally) in non-random patterns—is an open problem (Cromley et al., 2014). As many alternative places-as-aggregates can encompass the same entities at different scales, new heuristics are needed to construct optimal aggregates that meet user informational needs at a given scale, based on statistical measures of informational entropy. As I hope I have demonstrated, the uneasy relationship between space and place offers opportunities for unlocking new ways of searching the ocean of geo-information.

References

Agnew, J. (2011). Space and Place. In J. Agnew & D. Livingstone (Eds.), *Handbook of Geographical Knowledge* (pp. 316–330). London: Sage Publications.

Casey, E. S. (1997). *The fate of place: A philosophical history.* Berkeley, CA: University of California Press.

Cromley, R. G., Hanink, D. M., & Bentley, G. C. (2014). Geographically Weighted Colocation Quotients: Specification and Application. *The Professional Geographer* 66(1): 138–148.

Goodchild, M. F. (2011). *Formalizing Place in Geographic Information Systems*. In L. M. Burton, S. A. Matthews, M. Leung, S. P. Kemp, & D. T. Takeuchi (Eds.), *Communities, Neighborhoods, and Health* (pp. 21–33). New York: Springer.

Keßler, C., Janowicz, K., & Bishr, M. (2009). An Agenda for the Next Generation Gazetteer: Geographic Information Contribution and Retrieval. In *Proceedings of ACM GIS '09* (pp. 91–100). New York: ACM.

Purves, R., & Jones, C. (2011). Geographic information retrieval. *SIGSPATIAL Special* 3(2): 2–4.

Ballatore, A., Wilson, D. C., & Bertolotto, M. (2012). A holistic semantic similarity measure for viewports in interactive maps. In S. Di Martino, A. Peron, & T. Tezuka (Eds.), *Web and Wireless Geographical Information Systems* (pp. 151–166). Berlin: Springer.

Scheider, S., & Janowicz, K. (2014). Place reference systems. *Applied Ontology 9*: 97–127.

Thrift, N. (1999). Steps to an Ecology of Place. In D. Massey, J. Allen, & P. Sarre (Eds.), *Human Geography Today* (pp. 295–322). Cambridge, UK: Polity Press.

The VERP Explorer—
A Tool for Applying Recursion Plots
to the Eye-Movements of Visual-Cognitive Tasks

STUART K. CARD
Consulting Professor
Department of Computer Science
Stanford University
Email: scard@cs.stanford.edu

with

ÇAĞATAY DEMIRALP
Post Graduate Fellow
Department of Computer Science
Stanford University
Email: cagaytay@cs.stanford.edu

JESSE CIRIMELE
Graduate Student
Department of Computer Science
Stanford University
Email: cirimele@cs.stanford.edu

Designs in human-computer interaction (HCI) often involve trading between spatial and textual representations to achieve a nuance of representation that makes a task faster to execute, easier to learn, or less prone to error. Such designs can be very effective, but they can also be subtle, and it can be difficult to understand the mechanisms in play. Even generally successful interfaces can still hide bad combinations of interface, task, and context that could be improved were they identified.

One method of approaching this problem is to run chronometric experiments with contrasting conditions. Aside from being expensive for development work, this method is at such an aggregate level that it often does not provide much access or insight into the underlying mechanisms at work. Another method is cognitive simulation (Kieras, 2014). The intent is to specify the likely mechanisms at work and to validate them by their ability to predict chronometric or other data. The validated simulator can then be put to work on inferring other consequences of the design with some claim to knowing why. While this method has advantages, it is even more expensive and is most practical for large projects or projects close to an existing model that can provide a starting point.

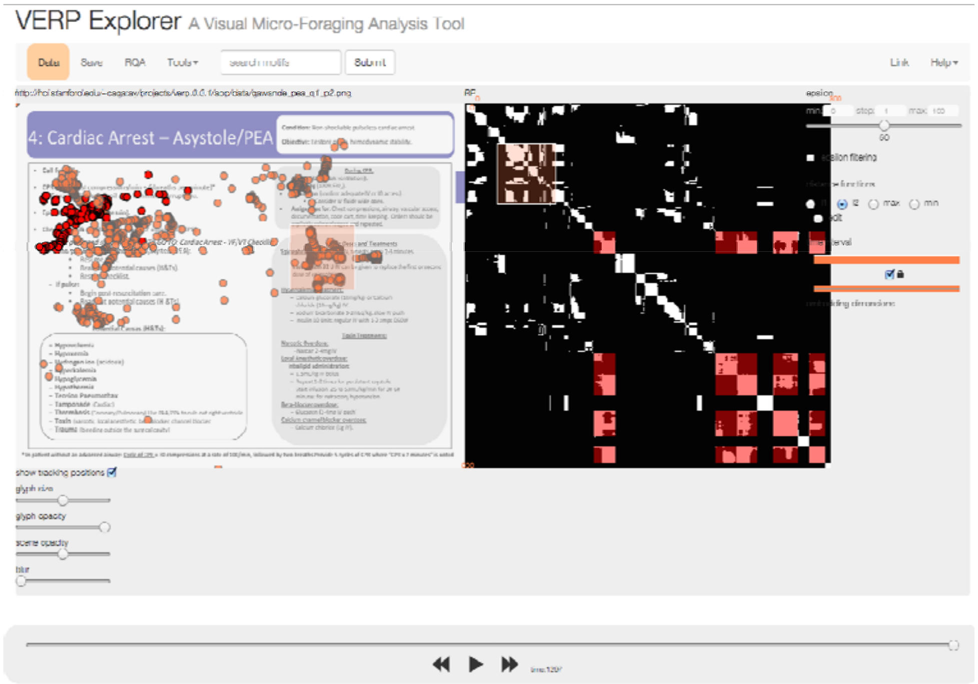

Figure 1. The VERP Explorer.

A third method is to construct a tool that makes the mechanisms at work visible by when applied to samples of user behavior. In this paper, we propose such a method and tool, *The VERP Explorer* (VERP stands for *Visualization of Eye-Movements based on Recurrence Plots*), an interactive visualization based recurrence plots. Eye-movement sequences are taken of users performing visual-cognitive tasks with the subject system. These are mapped into recurrence plot visualizations to highlight patterns of quasi-sequential behavior. In our system, these patterns are then back-mapped into—and overlaid on—the eye-movement scene to help characterize and provide insights into the behavior.

Recurrence plots are a type of non-linear analysis that has been used in the study of dynamical systems and other areas (Eichmann *et al.*, 1978; Marwan, 2008). Recently it has been applied to eye-movements (Anderson et al, 2013). Our tool extends and integrates eye-movement and recursion plot analysis into an interactive tool, simplifying exploratory analysis. Eye-movements can be thought of as a sequence of eye gaze positions f_i parameterized by time. To obtain the matrix [r_{ij}] that is the basis for a recurrence plot, we start with the first eye-position f_1 and compare it to all the other eye-positions in the sequence, including itself. If the distance $d(f_{ij})$ between the two compared eye positions is within some small distance $|\varepsilon|$, then we put a 1 at that position in the matrix, otherwise a 0.

$$r_{ij} = \begin{cases} 1, & d(f_{ij}) \leq |\varepsilon| \\ 0, & \text{otherwise} \end{cases}$$

We color white the cells whose value is 1 and black the others. Figure 1 shows an example of such a recurrence plot in the VERP Explorer. The data are a sample of eye-movements generating in our previous studies of the visual-cognitive task of emergency medical checklist use (Crimele et al., 2014; Wu et al., 2014). At the left in the figure is an image of the medical checklist, on which eye-movements from the eye-tracker have been superimposed. The doctor in this task was trying to use the checklist to answer the question, "What is the dose for Vasopressin?" On the right is the corresponding recurrence plot. As can be seen, the white areas are scattered in a disorganized fashion reflecting the disorganization of the search (Figure 1). Recurrence plots of other checklist designs are different, reflecting the visual-cognitive processes required to extract needed information. A small region of the recursion plot has been brushed (the nearly square brownish area). This "brushing" highlights some of the dots on the checklist eye-movement scene, causing them to turn red, indicating that the square-ish area on the recurrence refers to red text in that area. The VERP Explorer also allows us to go in the opposite direction. The area brushed on the eye-movement scene panel highlights the red areas on the recurrence plot, from which we can tell that the doctor actually looked at the area where the answer is, failed to see it, and came back for an intensive re-examination of that same area later.

Examining the recurrence graphs for different designs of the emergency checklists suggests the way in which the designs affect the details of the search. Many of the patterns can be broken down into "motifs" that signal certain types of behavior. These patterns can be quantified by recurrence quantification analysis (RQA) and thereby aggregated and compared. By using this tool, we have identified a pattern of search that looks a lot like that described by information foraging theory on a miniaturized scale. Doctors search for an information patch using searching saccades (explore), then read some information in the patch using fixations (exploit). If they fail to find what they want, they repeat the pattern. This pattern identified, of the doctor using multiple foraging cycles, suggests there is great room for improvement for this checklist and that the target is too similar to other targets or is surrounded by distracting elements, or is not unique, etc. With The VERP Explorer, we could quickly generate and test some alternatives. We believe VERP is a practical tool and we are making it into an application downloadable from the web.

References

Anderson, Nicola C, Bischof, Walter F., Laidlaw, Kaitlin, Risko, Evan F., and Kingstone, Alan (2013). Recurrence quantification analysis of eye movements. *Behavior Research* 45: 842–856.

Cirimele, Jesse, Wu, Leslie, Leach, Kirsten, Card, Stuart, Harrison, T. Kyle, Chu, Larry, Klemmer, Scott R. (2014). RapidRead: Step-At-A-Glance Crisis Checklists. *8th International Conference on Pervasive Computing Technologies for Healthcare* (Oldenburg, Germany, May 20–23, 2014).

Eckmann, J.P., Kmphorst, S. Oliffson, & Ruelle, D. (1987). Recurrence plots of dynamical systems. *Europhysics Letters*, 4(9): 973–977.

Kieras, David E. (2014). Towards accurate and practical predictive models of active-vision-based visual search. CHI 2014.

Marwan, N. (2008). A historical review of recurrence plots. *European Physics Journal Special Topics* 164: 3–12.

Pirolli, P. and Card, S. K. (1999). Information foraging. *Psychological Review 106*(4): 643–675.

Wu, Leslie, Cirimele, Jesse, Leach, Kirsten, Card, Stuart K., Chu, Larry, T Kyle Harrison, Klemmer, Scott R. (2014). Supporting crisis response with dynamic procedure aids. *ACM Conference on Designing Interactive Systems* (Vancouver, Canada, June 21–15).

Augmenting Intuitive Navigation at Local Scale

RUTH CONROY DALTON
Department of Architecture and the Built Environment
Northumbria University
Email: ruth.dalton@northumbria.ac.uk

With

JAKUB KRUKAR
Department of Architecture and the Built Environment
Northumbria University
Email: jakub.krukar@gmail.com

For long, internet search has been performed almost exclusively at a distance—from one's own desktop computer. When the search contains geographical information, it has been justifiable to provide the entire map of the environment. Such a map typically allows the user to integrate the represented information into his/her geographical survey knowledge. Most recently however, the increasing number of searches is performed from mobile devices, out in the world, and considers places located at closer proximity, within the spatial context of the searcher. And yet, in mobile-based navigation, the only substantial progress in relation to desktop-based (or traditional paper-based) representations is the presence of a "You-Are-Here" indicator. The traditional approach of providing the map of the entire relevant area in order to aid its integration into the searcher's existing survey knowledge thus presents an unmerited cognitive processing challenge. In addition, the sole action of looking at a map (not to mention its processing) seems superfluous. We argue that human-computer interaction for spatial searches should be (a) based on local, egocentric cues, intuitively preferred in human spatial behavior instead of global, allocentric, representations of the entire environment, and (b) confirming to the ideas of *HINTeractions* [1] and *calm computing* [2], where the output communicated by the system and hence intrusion is minimized [3].

The local context of a large proportion of spatial searches results in a typical user-case scenario in which the searched location must be navigated to by foot ("nearest pub," "open shop," "bus stop"). This task can be performed without referring to high-level representations of the environment. Knowing the user's position in the environment and his or hers orientation, it is possible to make use of the visual cues typically used in everyday navigation. These include, but are not limited to street width, or line-of-sight lengths of alternative spatial choices. Instead of suggesting the user to "turn right after 150 meters," pre-computing the map (i.e., using the geoinformatics approach) might allow the system to identify "that smaller street on the right," or "the main road in front." This natural language approach, even if consisting of fuzzy definitions, provides the opportunity for much more natural interaction [4–5].

While this method bears the risk of generating mistakes, it does not affect the precision of GPS tracking. As a result, despite the fact the user might take a wrong turn, the system has the capability to recalculate the route and repeatedly correcting any errors. According to the ideas of *calm computing*, this fact does not need to be indicated to the user, since no additional input is required to perform the task correctly. Not only the system should not alarm the user when it is not necessary, but it should avoid making explicit suggestions as long as it expects the user to perform well without cues. Numerous studies in Space Syntax and Spatial Cognition have demonstrated "default" patterns of spatial behavior during decision-making situations (such as the preference for paths providing longer line of sight). And yet, this knowledge base remains underused by navigational systems. Local spatial searches of instances to which navigation can be performed on foot exemplify a well-suited context for its application. When the system can reliably predict the default spatial decision most likely to be made by the user (and this prediction matches an optimal, or close-to-optimal solution), there is no need for providing explicit, complex output, like a map. Instead, the system can indicate that "everything is ok" and the user shall continue intuitively, according to one's own "default" strategy. Certainly, even during a short, on-foot navigation, there are critical decision points, at which making the correct decision is important, and no local cues suggest one. New wearable technologies make it possible to limit this output to the bare necessity, decreasing the cognitive effort required to process it.

We can consider three examples of distinct wearable technologies (augmented reality glass, smart watch, and audio-enabled earrings [6]) to demonstrate how the action of looking at a map can be replaced by unobtrusive visual, audio, or haptic *hints*. In each case, pre-processing of the required route takes place in the system, similarly to traditional geoinformatic methods. However, it is additionally enriched by the perceptual information available to the user from his/her current position.

In the augmented reality glass scenario, the user performs a search ("take me to the nearest ATM") from the middle of a market square. The glass responds by highlighting left side of the user's viewing field (or slightly darkening the rest of it). The user starts walking in this direction and comes across a junction where the market square ends and divides into two roads. Using Space Syntax analysis of the map of the area and visual image recognition, the glass recognizes that the left leg of the crossroad is perceived as the "main road," being much wider, more connected, and more densely populated with small businesses than the right leg. It therefore does not suggest the user any choice, but indicates (with a small green mark) that "everything is fine." The user, seeing the lack of explicit alerts, takes the preferred route and continues until the spatial configuration of the environment prompts the system to display another *hint.*

A similar scenario can be simulated with a smart watch, which can provide haptic feedback regarding the direction, but is limited in the number of choices it can suggest. Thus, small precise vibrations can serve as "left/right/straight on" *hints,* while more intense vibration alarms the user of a major mistake and suggests consulting the display for details.

Audio-enabled earrings can either use natural language commands, non-lingual audio signals, or a distortion to the music already being listened. A user turning to the wrong direction might hear a sudden decrease in volume, which immediately comes back to normal as the head orientation changes to correct for the decision error.

References

[1] G. Garcia-Perate, P. Agarwal, and D. Wilson, "HINTeractions: Facilitating Informal Knowledge Exchange in Physical and Social Space," in *Proceedings of the 3rd International Conference on Tangible and Embedded Interaction*, 2009, pp. 119–122.

[2] M. Weiser and J. S. Brown, "The coming age of calm technology," in *Beyond calculation*, Springer, 1997, pp. 75–85.

[3] W. Ju and L. Leifer, "The design of implicit interactions: Making interactive systems less obnoxious," *Des. Issues*, vol. 24, no. 3, pp. 72–84, 2008.

[4] R. Dale, S. Geldof, and J.-P. Prost, "Using Natural Language Generation in Automatic Route," *J. Res. Pract. Inf. Technol.*, vol. 37, no. 1, p. 89, 2005.

[5] H. Cuayáhuitl, N. Dethlefs, K.-F. Richter, T. Tenbrink, and J. Bateman, "A dialogue system for indoor wayfinding using text-based natural language," *Int. J. Comput. Linguist. Appl. ISSN 0976*, vol. 962, 2010.

[6] P. Drescher, D. Tan, W. Hutson, D. Berol, K. Merkher, R. Granovetter, A. Kraemer, K. Collins, and J. Rippie, "Group Report: Form Factors and Connectivity for Wearable Audio Devices," 2012.

Visual Search:
Guided Eye Movements for Foveated Sensory Systems

MIGUEL ECKSTEIN
Department of Psychological and Brain Sciences
University of California, Santa Barbara
Email: eckstein@psych.ucsb.edu

Find the parking spot. Find your backpack in the car's trunk. Find the blue garbage bin to throw the recyclable cup of coffee you were drinking in the car. Find the elevator. Find the button to operate the elevator. Find the key that opens the office. Find the keyhole. Find the outlet to plug in your laptop. Find the power button on the laptop. Find the icon for the email. Life is comprised of short visual searches. Each of these searches often involves moving the eyes to point the central area of the human retina (the fovea) to regions of interest in the scene to extract information related to the search. The fovea processes visual information with high spatial detail. Visual processing away from the fovea (peripheral processing) is mediated by a lower density of cone photoreceptors, higher convergence of cones onto retinal ganglion cells, and fewer associated neurons in primary visual cortex per millimeter of retina. Its consequence: Fine spatial discriminations are not possible with peripheral processing and thus humans often must make eye movements to utilize the fovea to explore a scene and search. But why have many animals evolved this varying resolution (foveated) visual system? Why not have, instead, a visual system that supports fine spatial detail across the entire visual field? The answer is likely related to the high metabolic cost of running a full high resolution system over the entire visual field. Over a ¼ of the brain is already dedicated to vision, and visual processing at full resolution across the entire visual field would increase that substantially. The density of cones in the fovea is approximately 20 times larger than at 10 degrees in to the periphery and 90 times at the far visual periphery (Curcio, Sloan, Kalina, and Hendrickson, 1990). The fovea which occupies 0.01 % of the retina utilizes approximately 10 % of the neuronal machinery in primary visual cortex (Azzopardi and Cowey, 1993). A high resolution processing system across the entire visual field matching the fovea's ratio of primary visual cortex (V1) neurons per mm of retina would result in roughly a one thousand size increase in the primary visual cortex. Instead, many animals have evolved a varying resolution visual system where a central area is given preferential processing and representation in the brain. However, a foveated visual system relies critically on the guidance of eye movements to efficiently explore

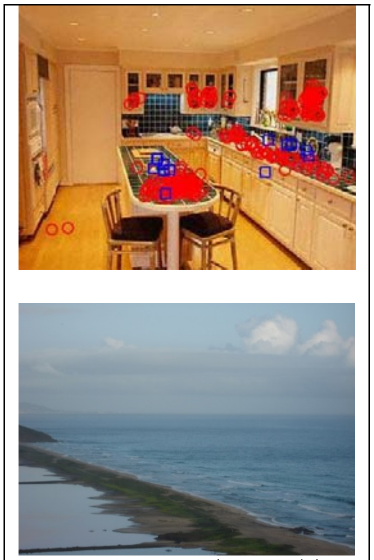
Figure 1. *Top:* endpoints (blue squares) of first eye movements of different subjects while searching for a coffee cup absent in the image. Red circles are selections by 100 observers of the location expected to contain a coffee cup; *Bottom:* Image of the ocean from the watchman's vantage point at a commonly used hill in Margarita, Island, Venezuela.

and extract information from scenes. Humans perform these eye movement searches effortlessly and automatically. The brain uses peripheral processing to extract critical information and guides the eyes across the scene. Eye movements are guided by information about the searched target including basic features including color, size, orientation and shape (Eckstein, Beutter, Pham, Shimozaki, and Stone, 2007; Findlay, 1997; Malcolm and Henderson, 2009). The brain is able to acquire information in the visual periphery to guide eye movements concurrent with analyses of information at the foveal region (Ludwig, Davies, and Eckstein, 2014). Yet often times a target can be small and difficult to detect in the visual periphery and the human brain must rely on other visual cues to guide eye movements. Objects in the visual environment are typically not randomly located: fruits tend to be on or under trees, plates are placed on tables, and chimneys on houses. Humans have a remarkable ability to learn statistical relationships among objects (Fiser and Aslin, 2002). The brain uses peripheral processing of a scene and learned relationships among objects or basic features to rapidly guide eye movements towards likely target locations (Figure 1; Eckstein, Drescher, and Shimozaki, 2006; Torralba, Oliva, Castelhano, and Henderson, 2006; Zelinsky, 2008). In searching for a coffee cup, the 1st eye movement (blue squares) is directed within 200-300 ms to countertops and locations that are in agreement with what subjects judge to be likely coffee cup locations (red circles: explicit reports about expected object locations in a scene).

The ability to use scene context to guide eye movements can be an indication of visual expertise in many perceptual tasks. For example, traditional fishing, still practiced in many places in the world, requires a watchman to sit for hours at the top of a hill and visually inspect the ocean surface for the presence of a particular type of fish (e.g., sardines). When the watchman detects the school, he will signal to his fellow fishermen to boat out to the ocean to surround the school with the appropriate nets. Because the presence of the school is often difficult to see on the ocean (Figure 1) with the visual periphery, the watchman typically relies on the presence of certain birds on the ocean surface to indicate the possible presence of the school. The birds' motion, easily visible in the periphery, guides the watchman's eye movements to the likely school location for further scrutiny with the fovea.

Figure 2. The lateral occipital complex (in red) is a brain region that represents the location in a scene humans expect to contain the target. Heat maps show the expected locations for a variety of targets across scenes collected from explicit judgments from subjects.

Where is the information guiding eye movements represented in the brain? In recent years, studies using functional magnetic resonance imaging (fMRI) have suggested that object selective cortex (OSC), previously thought to be only important in the recognition of objects, seems to play a critical role representing searched target in cluttered scenes (Peelen and Kastner, 2011, 2014). Activity in these areas, prior to the presentation of the visual stimuli, seems to mimic the visually evoked activity by the currently searched target. Prior to search for a person, OSC activity becomes more person-like and while searching for a car, it becomes more car-like. These findings suggest that OSC might represent a target template used to guide search. In addition, the lateral occipital complex (LOC), an area in OSC represents the likely location of a target object in the scene (Preston, Guo, Das, Giesbrecht, and Eckstein, 2013, Figure 2).

References:

Azzopardi, P., and Cowey, A. (1993). Preferential representation of the fovea in the primary visual cortex. *Nature*, *361*(6414): 719–721. doi:10.1038/361719a0

Curcio, C. A., Sloan, K. R., Kalina, R. E., and Hendrickson, A. E. (1990). Human photoreceptor topography. *The Journal of Comparative Neurology*, *292*(4): 497–523. doi:10.1002/cne.902920402

Eckstein, M. P., Beutter, B. R., Pham, B. T., Shimozaki, S. S., and Stone, L. S. (2007). Similar neural representations of the target for saccades and perception during search. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *27*(6): 1266–1270. doi:10.1523/JNEUROSCI.3975-06.2007

Eckstein, M. P., Drescher, B. A., and Shimozaki, S. S. (2006). Attentional cues in real scenes, saccadic targeting, and Bayesian priors. *Psychological Science: A Journal of the American Psychological Society / APS*, *17*(11): 973–980. doi:10.1111/j.1467-9280.2006.01815.x

Findlay, J. M. (1997). Saccade Target Selection During Visual Search. *Vision Research*, *37*(5): 617–631. doi:10.1016/S0042-6989(96)00218-0

Fiser, J., and Aslin, R. N. (2002). Statistical learning of new visual feature combinations by infants. *Proceedings of the National Academy of Sciences of the United States of America*, *99*(24): 15822–15826. doi:10.1073/pnas.232472899

Ludwig, C. J. H., Davies, J. R., and Eckstein, M. P. (2014). Foveal analysis and peripheral selection during active visual sampling. *Proceedings of the National Academy of Sciences*, 201313553. doi:10.1073/pnas.1313553111

Malcolm, G. L., and Henderson, J. M. (2009). The effects of target template specificity on visual search in real-world scenes: evidence from eye movements. *Journal of Vision*, *9*(11): 8.1–13. doi:10.1167/9.11.8

Peelen, M. V., and Kastner, S. (2011). A neural basis for real-world visual search in human occipitotemporal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *108*(29): 12125–12130. doi:10.1073/pnas.1101042108

Peelen, M. V., and Kastner, S. (2014). Attention in the real world: toward understanding its neural basis. *Trends in Cognitive Sciences*, *18*(5): 242–250. doi:10.1016/j.tics.2014.02.004

Preston, T. J., Guo, F., Das, K., Giesbrecht, B., and Eckstein, M. P. (2013). Neural representations of contextual guidance in visual search of real-world scenes. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *33*(18): 7846–7855. doi:10.1523/JNEUROSCI.5840-12.2013

Torralba, A., Oliva, A., Castelhano, M. S., and Henderson, J. M. (2006). Contextual guidance of eye movements and attention in real-world scenes: the role of global features in object search. *Psychological Review*, *113*(4): 766–786. doi:10.1037/0033-295X.113.4.766

Zelinsky, G. J. (2008). A theory of eye movements during target acquisition. *Psychological Review*, *115*(4): 787–835. doi:10.1037/a0013118

The Changing Problems, Databases, and Tools in Spatial Search

W. RANDOLPH FRANKLIN
Professor, Department of Electrical, Computing, and Systems Engineering
Rensselaer Polytechnic Institute
Troy, NY
Email: mail@wrfranklin.org

With

MARCUS V. A. ANDRADE
Federal University of Viçosa
Brazil
Email: marcus.ufv@gmail.com

This paper presents a view of the changing nature of search occasioned by larger databases and new software and hardware tools.

The problems that we want to solve are changing. Simple queries of a database to find those records r satisfying some predicate P(r) are easy. The more interesting queries involve coincidences. Find all pairs of records r,s such that P(r,s) holds. For example, search a set of aircraft track logs for near collisions.

Searching algorithms' data access patterns haven't changed much. Read operations are common, while write operations are rare, except when building a search structure in a preprocessing phase. Also, except while preprocessing, there is little interaction between the database elements.

What is new is that searching larger databases requires greater logical rigor during any computations. Unlikely errors from a sloppy algorithm design are now not so unlikely as the number of operations grows. An example would be a numerical roundoff error causing a topological error such as a point to be on the wrong side of a line.

The available hardware is changing. For all hardware, whether CPU or GPU, computation is becoming much cheaper than reading or writing. This applies to data in memory, in addition to data on secondary storage. Data compression becomes more useful. Transforming and computing with the data is often free.

Two complementary classes of hardware have recently become affordable: general multicore Intel CPUs and NVidia GPU accelerators. (We cite Intel and NVidia because they are leaders this year, but in a few years they might not even exist.) Problems that required servers are now doable on workstations.

Intel CPU workstations can now support very large main memory and multiple processors and cores. Even laptops can have 32GB of main memory. Imagine having a workstation with 1TB of main memory available for your database. Any record from that database is retrievable in roughly constant time (ignoring caches). That changes the appropriate search data structure.

Binary trees, quadtrees and octrees are silly. A B-tree with the root having 1,000,000 children, or a 1000x1000 uniform grid, is much faster. A massive hash table might be optimal.

That massive main memory is accessible by multiple CPU cores, each with multithreading. The author's 2012-vintage lab workstation is a dual 8-core Intel Xeon, with each core supporting 2 threads, for a total of 32 threads, and 128GB of main memory (expandable to 256GB). That is programmable by adding OpenMP directives to a properly structured C++ program. The directives instruct the system as to what operations can be performed in parallel. Loops where the different iterations do not depend on each other are common candidates. Reduction operations, such as summing over an array, also parallelize in OpenMP.

An example of a well parallelizable search operation is out Union3 algorithm and implementation. The input is a set of overlapping congruent cubes in E3. The output is the volume of their union. One search operation is to find all sets of three faces that intersect. Another is to determine which points from a set are not inside any of the input cubes. When processing 100,000,000 cubes, Union3 executes 10 times faster on 32 threads than on one. (In that case, of the 600M-choose-3 triples of faces, 395M intersected.)

Running a CUDA program on an NVidia GPU accelerator such as their newest model, the K40, is a more adventuresome solution, but with plusses and minuses. Massive parallelism is possible with the K40's 2880 CUDA cores. However: efficiently CUDA programming is complicated. The amount of available memory is very small—only 12GB. That memory has 3 levels of caching. The register level, about 100x faster than most of the K40's memory, has only 256KB per thread block. Data that has to be retrieved from the CPU's memory is even slower. The result is that one CUDA core is perhaps only 5% as fast as one Intel CPU core.

GPU programming also strongly prefers simple regular data structures, and code where each thread follows the same execution path on data that is adjacent to the data accessed by the previous thread. Complicated algorithms and structures with recursion and pointers are deprecated. Elevation DEMS are better than TINs. This theme of the optimal solution being simple has occurred many times before, e.g., with hash functions and page replacement algorithms.

Because the physics of the current technology limits processor speeds to about what they are now, whether with CPUs or with GPUs, parallel programming is necessary for searching large databases.

(Switching a transistor is like charging a capacitor. Reducing the switching time requires increasing the voltage or making the transistor smaller. The former is infeasible; instead voltages get smaller to reduce power. The latter (counterintuitively) increases the transistor's resistance, which also requires more power. That runs up against the limits of how quickly heat can be removed from a small circuit.)

The available software is changing. Exactness may seem to be a pointless luxury in GIS applications, since GIS data is by its nature approximate, and approximate results are sufficient for all practical purposes. However, most geometric algorithms used in GIS, such as point location and map overlay, become much more complex and prone to failure if their elementary operations are subject to rounding errors, no matter how small. Consider for example a distributed application that cuts a map into smaller submaps, handles each piece to a separate processor, and combines the partial results into a single map. If the cutting step is exact, the final step needs only to identify common boundary edges between the partial results, and remove them. The task becomes much harder if the cutting step is affected by rounding errors: the partial results may overlap, or may be separated by gaps. The pasting operation is then almost impossible to specify, let alone to implement.

The well-founded solution to roundoff errors in geometric computing is to compute in a number domain that does not suffer roundoff. That would be the rational numbers. Each number is composed of a numerator and a denominator, each of which is an integer with as many digits as necessary (implemented as an array of C++ ints). gmp++ is one library that implements rational computation. Computing with rationals is slower than computing with floats; however see the earlier comments about computation cost.

Searching and working with curves as curves, instead of as a polyline approximation, is starting to look like perhaps becoming feasible. CGAL (Computational Geometry Algorithms Library) is an example of a very large package for more complete geometric computing. The intersection point of two curves can be determined as the root of a polynomial equation. (Just as common variables can be eliminated from a set of linear equations, so can a set of polynomial equations be solved by techniques such as resultants.) The problem is that this polynomial is of a very high degree. A bicubic parametric patch, the lowest degree that allows matching curvature when two patches join, has degree 18 when considered as an implicit function. Intersecting two produces a curve of degree 324. This is an area of continued research.

These powerful software tools are now feasible only because of the newly affordable powerful hardware. They come at the right time to facilitate higher-order searches in the new massive databases.

Search in Spatially Structured Worlds

CHRISTIAN FREKSA
Bremen Spatial Cognition Center
University of Bremen
Email: freksa@uni-bremen.de

The notion of spatial search is used in at least three different senses: (1) search in real (vista, environmental, or geographic) space; (2) search in virtual space; and (3) search in non-spatial domains using a spatial metaphor for navigation. Important distinctions of real spaces in comparison to other spaces are: their inherent restriction to three, two, or one dimension; the presence of at most one physical entity on any location; and a strict ordering of entities on any of the dimensions (with periodic orderings on circular dimensions). These properties can be exploited for effective and efficient search procedures. Unless otherwise noted, I will focus on search in real space.

In comparison to abstract spaces, size-restricted real spaces contain a finite number of physical entities; but although at any given location we only can have a single physical entity, this entity may belong to a multitude of conceptual entities; these may be in a hierarchical or heterarchical relationship to one another. These properties of entities in real space and their conceptualizations are responsible for interesting features related to search.

In abstract spaces, we may have an arbitrary number of dimensions such that each entity can be a neighbor of each other entity; this would allow us to directly reach any entity from any other entity; the problem with this very general framework is that it does not scale well with the number of objects involved: each new object requires a link to each existing object; this constitutes an exponential growth with respect to the number of entities, and consequently results in expensive computational procedures.

If we restrict the dimensionality of space to those of real spaces, we also restrict the number of spatially neighboring physical entities; but we do not seriously restrict the places we can visit: we first visit a neighbor; from there we can visit a neighbors' neighbor, etc. If all entities are connected in a single network, we can reach any place. The growth in complexity is only linear in the number of entities involved. Reaching nearby places is cheaper than reaching distant places; this is exactly what we want in a benign world in which Tobler's First Law of Geography (*"Everything is related to everything else, but nearby things are more related than distant things")* applies (Tobler 1970). The property described here applies on any level of spatial resolution.

Spatial neighborhood and spatial movement induce conceptual neighborhood of spatial relations (Freksa 1991). Together, they result in a very interesting and useful property that allow for the conceptualization of spatial entities on various levels of spatial and conceptual resolution. For example, we can conceptualize an entity at a given spatial location as a house, as part of a settlement, or as a room (in a house); this is a side-effect of spatial coherence that is expressed in Tobler's First Law.

The spatial coherence between fine-grained and coarse-grained spatial entities and their corresponding conceptualizations enable very useful approaches to spatial search; in particular, we can move to a coarser level of resolution (zoom out) when we search for a certain entity; this will reduce the size of the search space while at the same time this will widen the scope to closely related items that may not fully agree with the fine specification the search has started with. At a later stage of search, we can refine resolution (zoom in) again in order to select the most suitable candidate from the answer set. In human cognition, such coarsening and refinement operations appear to constitute important mechanisms for efficient and effective search.

Here are some issues to discuss in connection with spatial structure and spatial integrity:

Conceptual dimensions of spatial structure:
- horizontal (near – far) vs. vertical (fine – coarse)
- search metaphors: fine to coarse and back vs. coarse to fine and back
- abstraction – concretion operations
- approximation – precesiation operations

Representational / cognitive economy
- why do we coarsen knowledge obtained in fine structures when we require precise answers?
- why do we refine knowledge when coarse answers are sufficient?

The role of qualitative spatial relations in search

Exploitation of redundancies

Spatial context and context-based affordances

Virtual vs. real search spaces

The role of spatial reference systems for search
- location vs. content search

Knowledge in the world vs. knowledge in the head
- quantitative / qualitative interaction (e.g., wayfinding in the real world)
- spatial and conceptual neighborhood (Tobler's 1st law of Geography)

Trade-offs
- space – time – accuracy
- one or all search

Affordance- vs. constraint-based search
- forward – backward search
- incomplete and / or inaccessible knowledge

References

Freksa C., Conceptual neighborhood and its role in temporal and spatial reasoning, in Singh M., Travé-Massuyès L., eds, *Decision Support Systems and Qualitative Reasoning*, 181–187, North-Holland, Amsterdam 1991.

Tobler W. (1970) A computer movie simulating urban growth in the Detroit region. Econ Geogr 46(2): 234–240.

Guided Spatial Search in Digital Maps

WAI-TAT FU
Associate Professor, Department of Computer Science
University of Illinois, Urbana-Champagne
Email: wfu@illinois.edu

In the field of cognitive science, search is often considered a fundamental process for a wide range of cognitive activities. The reason why search is fundamental is partly because when cognition needs to decide on the next operations to achieve its goals, it often does not have perfect knowledge to make that decision. Instead, multiple *cognitive search* processes are needed to determine the available operations and to evaluate which operations should be chosen to achieve its goals. For example, in problem solving, *state-space search* involves the exploration of the external environment to determine what actions are available, and *knowledge search* involves the processing of internal knowledge to evaluate the actions to determine the best next action. These search processes allow the problem solver to traverse the problem space *intelligently* to reach the goal (i.e., solve the problem), and are therefore fundamental to understanding the cognitive computations involved in various problem-solving situations.

In general, maps are abstract external representations of structures in the spatial environment. Like most external representations, they contain symbol structures that can be perceived and recognized, and their relations can be inferred and conceptualized such that knowledge about the represented entities can be acquired and integrated to *make sense* of the information. For example, a *location* on a map becomes a *place* or a *meaningful spatial object* when it has acquired meaning that allows someone to infer the attributes of the place (e.g., a park, a train station, or a University). However, maps are also constructed specifically to convey spatial relations of locations or spatial objects, and these relations are represented in ways such that they are maintained at many granularity levels. Most importantly, spatial relations at different granularity levels often allow users to extract and integrate knowledge about the spatial environment, and to process the spatial relations in ways that help them to accomplish their goals.

This position paper aims to discuss to what extent one can characterize cognitive processes of a person when engaged in spatial search using digital maps. In digital maps, actions such as *zooming* and *panning* allow users to search for locations or meaningful spatial objects as represented on the map. One important characteristic of digital maps is that, given the limited screen size, there is an inherent trade-off between *spatial context* and *details*—while one can zoom out to derive the relative locations of multiple objects on the map, some specific objects may not be shown or their locations are often imprecise; but when one zooms in for details of location names, roads, or intersection, one will lose the broader spatial context of where the map is representing relative to other locations not shown on the map. Sequences of zooming and panning actions are therefore needed to maintain continuity of the representation of spatial objects, which often impose unique challenges to spatial search.

To the best of our knowledge, there has been a lack of empirical studies on how people perform spatial search using digital maps, let alone developing computational cognitive models for the activity. To better understand spatial search using digital maps, we designed an experiment that required a pair of participants to collaboratively search for a sequence of locations distributed across multiple zoom levels and locations, and eventually identified a region on the map. The goal is to understand (1) how people communicate locations on digital maps at different zoom levels, (2) how people combine zoom and pan actions to handle the trade-off between context and details during spatial search, and (3) how people develop collaborative strategies incrementally when performing spatial search using digital maps.

In the spatial search task, two participants were separated in different rooms. They could not see each other but they could communicate over an Internet phone, and each of them sat in front of a desktop computer and interacted with the experimental interface that showed the digital maps. The digital maps were created using the Google Map API. Each participant would be given a different set of location "pins" distributed at multiple zoom levels of the digital map of a large U.S. city (either Chicago, San Francisco, Seattle, Boston, and Pittsburgh). The pair of participants took turns to be the director and the matcher. The director would first locate one of the location pin, and guide the matcher though the Internet phone to search for the pin on the matcher's map. Once the matcher believed that s/he found the pin, the matcher would place the pin at a particular zoom level (which might not be the same zoom level as presented on the director's map). After one pin was placed, the participants switched roles until all pins were placed. Participants were then told that to infer and select the regions enclosed by all the pins on the map.

I will discuss some interesting findings from this study, and the extent to which the results shed light on the cognitive aspects of how people perform guided spatial search using digital maps; and whether it has implications to design of spatial information systems. I will also discuss the development of a computational cognitive model that characterizes the collaborative spatial search task, and how and whether certain processes are similar or different from models of information search. It is expected that the results will provide a better understanding of the construction and changes of cognitive representation of locations during the search process. Results may also shed light on design of digital maps or relevant tools that facilitate spatial search.

Code, Content, and Control in Spatial Search[1]

MARK GRAHAM
Associate Professor, Oxford Internet Institute
Research Fellow, Green Templeton College
Associate, School of Geography and the Environment, University of Oxford
Email: mark.graham@oii.ox.ac.uk

Organizations that facilitate practices of spatial search are playing an ever more central role in the lives of billions of people.[2] This is because they are making editorial, economic, and ethical decisions that shape our lives (Kitchin and Dodge 2011; Graham et. al. 2013). Organizations like Google and Apple, and platforms like Wikipedia and Facebook are arbiters of not just what we see and read, but what we know about our world, and how we navigate through our world.

Because of that, and as they become ever more integral to our lives, it is instructive to turn to some provocations of the late Tony Benn, a British Labor party politician. He famously had a set of five questions that he said that we should always ask any powerful person: *"What power have you got? Where did you get it from? In whose interests do you exercise it? To whom are you accountable? And how can we get rid of you?"* This paper is organized around those questions, directing them to dominant practices of spatial search:

What power have you got?

Entities that facilitate practices of search are curators and mediators of information[3] (Graham, Schroeder, and Taylor 2014). Searches are usually many-to-one interactions that can lead us to different: *representations or truths*; *voices or controllers of information; places* (Graham et al. 2013).

Where did you get it from?

Entities that facilitate practices of search are able to exert power in five primary ways:

Through algorithms. Much of the curation and mediating is carried out through opaque algorithms that are able to respond to a query with targeted lists, layers, or maps of information. These algorithms are often opaque to end users and undoubtedly privilege some content over others. Google, for instance, has long since abandoned a reliance on its PageRank system to determine what content to make visible or invisible in response to any particular query. Instead, it employs a much complex array of code that accounts for a diverse range of locally, temporally, and individually contextual information. While this algorithmic targeting of information undoubtedly has benefits to end-users (e.g., a user in Delhi receiving locally-relevant content if searching for "what time is it?"), it can also create "filter bubbles" that can prevent users from encountering representations of the world that clash with their own perspectives (Graham and Zook 2013). A Chinese and an American search engine user will receive very different imagery if conducting a search for "Tiananmen Square."

Through data/system architectures. The design and structuring of the tools that we use to conduct searches can exert important influences on how we obtain spatial knowledge. One well-referenced example is the fact that only a small percentage of search engine users ever click

---

[1] Because of a three-page word-limit, this is an abridged version of a longer paper available at geospace.co.uk

[2] Here I consider two types of spatial search: (1) *explicit spatial search*: in which the object of a search is a place; (2) *implicit spatial search*: in which the object of a search is spatial information (often about a place).

[3] One only needs to look at the size of the search engine optimization market (an industry valued in tens of billions of dollars) to see the economic value that has been placed on attempting to influence those mediations.

through to a second page of search results. This means that for all but a minority of people, a search engine's decision to place something on the first page is a decision to effectively make that thing visible or invisible.

Through data presences and absences. It is crucial to remember that even in our era of "big data," there remain stark differences in the amount of digital content created and indexed about different places and processes (Graham 2014). The "data shadows" of global cities and high-income countries are thick and dense: with many indexed photographs, descriptions, tags, encyclopedia entries, and other geographically grounded content. Yet much of the world remains characterized by thin and sparse "data shadows"—businesses with no digital presence, potential tourist attractions and sites of interest invisible on maps, and important narratives of places either missing or created by people and organizations from the world's informational cores.

Through the systems of governance. The ways that platforms, organizations, and systems of information are governed have a tremendous effect on practices of spatial search. The move toward structured data and the semantic web is, for instance, allowing for the creation of shared meaning across digital contexts (i.e., allowing content to be more easily separated from its containers and contexts). While shared meaning is useful and productive, it can also work to eliminate shared meaning. Returning to the earlier example, while different populations of Jerusalem exist in different language versions of Wikipedia, search engines harvest answers to queries such as "what is the population of Jerusalem" from structured databases like Freebase or Wikidata. These central databases necessitate agreement where disagreement exists (which rarely favors people or groups with minority perspectives).

Furthermore, large entities that facilitate search are able to exert power in one additional way: *through centrality in an attention economy*. Network effects exist that make it difficult to dislodge dominant platforms and services. For instance, if a search engine has a dominant share of human attention, it is better able to test, target, and respond to human behavior. This allows it to retain users, which creates at virtuous cycle for the company. The same can be said for spatial searches conducted via social networks or user-generated encyclopedias.

Practices that are mediated by spatial search are usually influenced by the confluence of many (if not all) of these factors. Microsoft, for instance, registered a patent dubbed by commentators as its "avoid ghetto" feature (Tashev et al. 2007). The patent allows users to be routed from one place to another based not just on time or cost distances, but also socioeconomic variables (like crime statistics). Poor neighborhoods with high crime statistics would then become digitally blacklisted. In this example, it is the confluence of algorithms (the routing system), data and system architectures (e.g., the design of a user-interface to encourage users to use the feature), data presences and absences (e.g., the presence or lack of data that allows a measure of risk to be calculated), systems of governance (choices made by engineers or executives to include that particular feature), and centrality (the propensity of users to avoid trying new and unknown services) that would allow people, their attention, and their economic activity to be channeled in some directions over others.

In whose interests do you exercise it? To whom are you accountable?

The full version of the paper (available at geospace.co.uk) addresses this question in much more detail. But the short answer is that almost all large organizations involved in spatial search are accountable to only their owners and shareholders.

How can we get rid of you?

Practices of search are always imbued with knowledge politics. Because there is rarely a singular correct answer or response to any query, mediators of search are always making inherently powerful decisions to make knowledge visible or invisible, to promote or demote certain perspectives, and to represent or hide particular places and processes. Furthermore, spatial search, even more than other types of search, highlights some of the power that information mediators have. The opportunity cost of performing a spatial search and then using those results to navigate to more than one place is comparatively high. This gives mediators of spatial search a tremendous amount of power to influence not just what we know, but also what we do and where we go. Because mediators of spatial search are rarely managed with the wellbeing of their users or the wellbeing of society as a priority, we should be asking what alternatives might look like.

Despite the litany of abuses perpetuated by the nation-state with respect to the mediation of information (e.g., excessive censorship or intrusions of privacy), it remains the only mechanism of governance that we have in which each person has one (and only one) vote. As such, it remains a potentially powerful force to ensure that minimum standards or best practices are adhered to. We might, for instance, imagine anti-discrimination laws being more cleverly applied and enforced in the contexts of spatial search. We might also see a need for explicit manipulations of the mediation process (e.g., paid advertising) being more clearly labeled as such. The state's role in education could also be harnessed to make "critical code studies" and "critical data studies" a much more prominent component of universal education. However, the risks of heavy state involvement are significant, and it remains that much can be done through horizontal rather than top-down initiatives.

Search engines survive in an attention economy. Thus, if any organized re-direction of attention occurs (such as through boycotts or the construction of alternatives), the organizers of that change have the power to enact important change in the world: either by shifting attention to other mediators or by encouraging existing mediators to change their practices. Some priorities are: the need for more serendipity to avoid people being trapped in informational filter bubbles, the need to make not just the "known knowns" visible as search results, but also the "known unknowns," the need for much more visible debate about who benefits and who loses out from any particular act of mediation, and the need for more transparency about some of the key logics used to make some information, processes, and places visible, and others invisible.

Benn's five questions offer a useful lens to guide inquiry into who benefits and who doesn't from particular modes of the capture, storage, flows, and presentations of information. But they also lead us to create alternate visions and ask how the regulation, manipulation, and mediation of information might be carried out in very different ways. There are many indications that currently dominant entities that facilitate practices of search are amplifying rather than circumventing the power of the already most visible, most wealthy, and most dominant; that predominant practices of search do nothing to reconfigure virtuous and vicious cycles. Benn's provocations, as well as conversations such as the one that this paper is a part of, give us an opportunity to hold entities that facilitate practices of search accountable for the power that they wield and re-imagine what more ethical or socially just practices might look like.

Geographic Search

KARL GROSSNER
Digital Humanities Research Developer
Stanford University Libraries
Stanford University
Email: karlg@stanford.edu

Spatial searches often require or benefit from contextualizing temporal and thematic parameters. The research agenda undertaken in this meeting should therefore aim to better theorize the much broader *geographic* search, and to develop pragmatic formalisms and technologies that integrate space, time and theme. This perspective is of course not new; it has been explicitly framed in Brian Berry's *geographic matrix* (1964) and by many others since, notably including May Yuan, who discussed search specifically (1999).

Many scientific inquiries consider spatial characteristics somewhat in isolation, by holding time and one or several themes constant, but an integrative view of the three dimensions is increasingly essential. Several "spatial turn(s)" are now ongoing across the academy,[1] but particularly in the humanities, examinations of space are rarely purely spatial in a geometric or topological sense. Rather, they are about place—spaces as named, constructed, and experienced by people. As such they are inherently geographical, even if there is some reticence to use that word. The potential benefit of place-based, or "placial" indexing of information becomes increasingly apparent in the humanities and social sciences, for exploratory research, for analysis, and for publication of *interactive scholarly works* (Meeks & Grossner 2012).

The Association of American Geographers (AAG) has recognized this trend for some time, evidenced in recent annual meetings having numerous extended tracks on the topic of gazetteers, and in the recently announced creation of a new journal, *GeoHumanities*. I was pleased to co-found in 2013 the *GeoHumanities SIG*, a special interest group within the Association of Digital Humanities Organizations (http://geohumanities.org). One focus of the SIG will be to help make digital humanists now struggling with geographic representation and computing issues aware of relevant work in GIScience. Conversely, GIScience will find many fascinating challenges it can contribute to solving. The first issue taken up by the SIG at this year's DH2014 conference in Lausanne was historical gazetteers.

There is also currently under way an NEH-funded global collaborative effort to develop a *world historical gazetteer*. At a formative meeting in September, 2014 I presented some of the following recent work I've undertaken on the necessary and often-skirted temporal dimension of historical gazetteers.

Historical time for indices and analysis

Spatial questions in the humanities are very often historical, making the temporal dimension at least co-equal with the spatial in those cases. For historians, existing temporal indexing schemes and computation methods are woefully inadequate. A large proportion of historical time

---

[1] An excellent series of essays by historian Jo Guldi: http://spatial.scholarslab.org/spatial-turn/

references are vague, probabilistic, or otherwise uncertain. I am currently developing with colleague Elijah Meeks a data format, graphical timeline layout, and computing functions for complex and uncertain historical time, called Topotime (http://dh.stanford.edu/topotime). We are actively pursuing its connection with both PeriodO, a promising initiative to build an authoritative time period gazetteer (http://perio.do/), and Pelagios (http://pelagios-project.blogspot.com), an important historical gazetteer integration project well under way.

Place and Period in Settings

Answers to "where?" questions can always include "when" (and often must); and so we can say that all phenomena occur within a *setting* that is equally spatial and temporal, whether or not both kinds of data are available or of interest in a particular study. I have recently begun work with Krzysztof Janowicz and Carsten Keßler on an ontology design pattern for *Setting*. Our pattern follows on from the GEM model of Worboys and Hornsby (2004) in some respects. It is motivated in part by the circumstance in historical studies that some entities are equally places and periods. The Iron Age is not one timespan everywhere; there is an Iron Age Britain, Iron Age Levant and so forth.

Our pattern describes a given place or period as being *scopedBy* a Setting, in turn comprised of a *SpatialScope* and a *TemporalScope*, each of which describes an extent in familiar terms of timespans and footprints. The pattern corresponds at least superficially with the Dublin Core *coverage*. One obvious track for future work will be consideration of whether the *ThematicScope* for a Setting should be included.

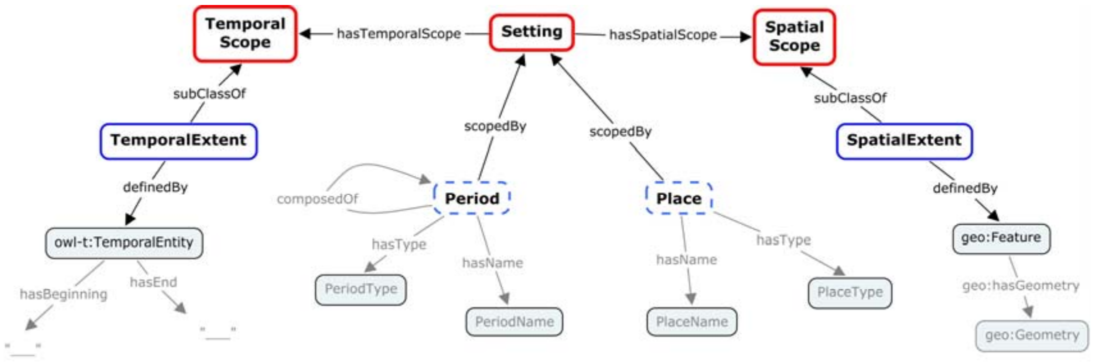

Theme

The thematic dimension of geographic search comes into play in several ways, including humanists' requirement for *multivocality* in knowledge representation. Multiple attributed accounts of the same phenomena must co-exist in the same system—the *open-world assumption*. An essential component of description is classification, and particularly in searches across disparate data sets, methods for both implementing formal semantics and/or mapping between disparate formal or informal vocabularies are central to the "thematic" aspect of geographic search. How can we unify *semantic reference systems* (Kuhn 2003) with those for space and time?

Works cited

Berry, B. J. (1964). Approaches to regional analysis: a synthesis. *Annals of the Association of American Geographers 54*(1): 2–11.

Grossner, K., Janowicz, K., and Keßler, C. (in press) Place, Period, and Setting for Linked Data Gazetteers. In R. Mostern, H. Southall, M.L. Berman (Eds.) *Placing Names: Enriching and Integrating Gazetteers*. Bloomington, IN: Indiana University Press.

Kuhn, W. (2003). Semantic reference systems. *International Journal of Geographical Information Science 17*(5): 405-409.

Meeks, E., and Grossner, K. (2012). Modeling Networks and Scholarship with ORBIS. *Journal of Digital Humanities 1(3) (Summer 2012).*

Worboys, M., & Hornsby, K. (2004). From objects to events: GEM, the geospatial event model. In *Geographic Information Science* (pp. 327–343). Springer Berlin Heidelberg.

Yuan, M. (1999). Use of a Three-Domain Representation to Enhance GIS Support for Complex Spatiotemporal Queries. *Transactions in GIS 3*(2): 137–159.

Use Cases and Personas for Spatial Search

DARREN HARDY
GIS Software Engineer
Stanford University
Email: drh@stanford.edu

With

JACK REED
Geospatial Web Engineer
Stanford University
Email: pjreed@stanford.edu

In this paper, we describe the use cases and personas we developed for our GeoBlacklight discovery services, an open source, multi-institutional software project started at Stanford University Libraries (http://geoblacklight.github.io). It provides discovery services across a federated multi-institutional repository of geospatial resources, and is part of a larger effort to preserve and provide access to geospatial data. This effort is an open collaborative project aiming to build off of the successes of the Blacklight (http://projectblacklight.org) and the multi-institutional OpenGeoportal federated metadata sharing communities (http://opengeoportal.org).

Our aim is build an application that enables discovery with an emphasis on user experience, integrates seamlessly with other web mapping tools, and streamlines the use and organization of geospatial data. Specifically, we focus on discovery and user experience, leaving analysis to other tools, on growing a community of open source development and metadata sharing partners, enabling discovery across institutions. Traditionally, geolibraries have focused on spatial search and organization, but not as much on federation and modern search semantics like query intent processing and faceting. Emergent geoportals, for example, have suffered from usability issues, lack of federation, and integration with other applications. GeoBlacklight's modularity as a Ruby on Rails engine lends itself to be deployed in a variety of Blacklight applications and contexts. It uses a metadata schema that enables specific geospatial discovery use cases for search, view, and curation functionality.

Use cases. We focus on discovery-oriented use cases for search, view, and curation, supported by six personas in a research library context (see Hardy & Durante (2014) for further details and diagrams). The search and view use cases include three modalities of text search, faceted refinement, and spatial search, and two actors: a casual user and a geoportal user. In general, current features in geoportal and web mapping applications are map and text views of results, clustered results, faceted search, text search, spatial search, preview layer, related items, groupings, and suggested search. Hardy & Duarte (2014) discuss how "search users either use web search engines or geoportal search directly. This is a very important design point as web search is so heavily engrained into users' workflows for discovery of resources in general." The primary search use cases are web search, geoportal text search, geoportal faceted refinement, and geoportal spatial search. For the view landing pages, the use cases are view dataset, view metadata, visualize layer, and download layer.

On the curation side, there are two actors: a curator and an administrator. The curation use cases cover how the geoportal application builds and manages its federated multi-institution repository of holdings. The primary use cases are to import and export holdings, curate holdings, monitor operations, and acquisitions via purchasing policies, archiving public layers, and self-deposit.

Personas. We divided our personas into Application End Users and Project Stakeholders and below we list their primary goals. For Application End Users, we have a Professor, PhD Candidate, and Student:

(a) Professor—experienced scholar, not a GIS user, looking for specific data:

- Quickly find historic maps for my area of study
- Use a tool that thinks like I do to find the information I need,
- Point my students to a great campus resource for data and historic maps.

- (b) PhD Candidate—power user, experience professional, focused on getting the job done:
  - Quickly evaluate potential data sources
  - Download and save discovered layers
  - Discover hard to find localized data sources

(c) Student—new to GIS and research, tech savvy, eager to learn:

- Finding GIS data to get started seems daunting
- Hard to determine which data are reliable
- All of these GIS tools seem to have been designed 10 years ago

and for Project Stakeholders, we have a Librarian, Lab Manager, and Web Engineer:

(a) Librarian—passionate about geospatial and access, wants to provide more patron services:

- Provide visual catalog to libraries' GIS data and maps
- Give patrons access to other institutions shared GIS data and metadata
- Understand how and what users are downloading/using the resources we provide

(b) GIS Instructor and Lab Manager—experienced professional, wears many hats, helping and training many:

- Easily share GIS data sources with students
- Provide instructions and demonstrate discovery of GIS data
- Have one place for students to go to obtain GIS data

(c) Librarian Web Engineeer—code contributor, open source advocate:

- Install and customize an instance of GeoBlacklight for university's users.
- Be able to contribute code to the GeoBlacklight project.

- Share metadata records with community.

These use cases and personas have guided our development to focus much more on user experience in light of modern search semantics, where spatial search is an enhancement rather than focus. Thus far GIS has made significant strides on the technical details behind spatial indexing and "GIS-styled" user interfaces and should now a focus on modern search semantics which are aided by spatial search and federation.

References

Hardy, D. and K. Durante. 2014. A Metadata Schema for Geospatial Resource Discovery Use Cases. *Code4lib Journal* (25). From http://journal.code4lib.org/articles/9710

Hearst MA. 2009. Search User Interfaces. Cambridge: Cambridge University Press.

Cognitive Perspectives on Spatial Search

MARY HEGARTY
Associate Director, Center for Spatial Studies
Professor, Department of Psychological and Brain Sciences
University of California, Santa Barbara
Email: hegarty@spatial.ucsb.edu

I do not see search as a unified topic in cognitive psychology, but issues concerning search are central to many theories and phenomena in cognitive psychology. Here I discuss some different ways in which perceptual and cognitive psychologists have considered search. I give examples representing three main topics, search of real spaces (e.g., the visible environment), search of information spaces in the mind (e.g., memory) and search of information spaces outside the mind (e.g., the world wide web).

Search of real spaces:

Visual search refers to our ability to search the immediately visible world to find some object or feature. An everyday example is searching for your keys on a cluttered desk. Classic research on visual search has focused on simple artificial stimuli such as searching for a red X among a pattern of letters of different colors, or detecting a pattern that is barely visible. One of the classic results in the literature, is that spatial cueing (i.e., receiving an explicit cue of where the target will likely be) increases the detectability of stimuli and speeds up search. One current direction in visual search is to study search in more naturalistic stimuli such as real scenes, and how people infer and use the statistical properties of the environment to guide search. For example people use information about where objects are likely to occur to detect objects more quickly in expected than in unexpected locations (Eckstein, Drescher, and Shimozaki, 2006).

Another issue is how animals and humans search for hidden objects in an environment, specifically, what spatial cues and other cues are used in this search context. In one classic set of studies (Cheng, 1986) animals were first familiarized with the location of food in a rectangular environment that also had featural cues (e.g., colored or textured walls) that could be used to locate the food. After disorientation, they had to search for the food (now hidden) in the same environment. A striking result is that animals searched primarily based on the geometry of the space and ignored distinctive cues such as the brightness and texture of the walls. Similar effects were found when young children search for hidden objects, although older children combine featural and geometric cues. While the interpretation of these results is controversial (Cheng and Newcombe, 2005), they suggest that environmental geometry, a spatial property, may be more fundamental than featural cues (e.g., distinctive visual and olfactory cues) in directing search.

Search of Information Spaces in the Mind:

Human memory can be conceptualized as an information space, and retrieval of memories can be thought of as search of this space. A basic phenomenon in memory is context dependent memory, the effect that memories are more easily retrieved when the context at encoding and retrieval is the same and spatial location is a powerful contextual cue. Thus memories are more easily retrieved in the same location in which they were originally stored, that is, the location in which a memory was encoded is part of the memory and can serve as a retrieval cue (e.g., Godden and Baddeley, 1975). The Method of Loci, an ancient mnemonic device used by orators to remember the points in a speech by associating them with places in a known environment, capitalizes on space as an organizer of memory.

A classic theory of problem solving (Newell and Simon, 1972) provides a different example of how psychologists have thought about search in an information space. According to this theory, problem solving activities, such as playing chess or solving the Tower of Hanoi problem, can be thought of as search in a "problem space", that is, the space of all states of knowledge in solving a problem, connected by operators for changing one state into another and constraints on applying these operators. According to this theory, problem solving involves search for an efficient path from the given state to the goal state of a problem. Newell and Simon proposed that rather than searching the problem space exhaustively, people rely on simplifying heuristics to narrow down the search process, an example of bounded rationality. This idea is central to current theories of ecological rationality (e.g., Todd and Gigernzer, 2007) which will be represented at the meeting.

Search of Information Spaces Outside the Mind

Finally, cognitive scientists have studied how people search information spaces that occur outside the mind, such as the world-wide web. One relevant line of research is information visualization (Card, Mackinlay, and Scheiderman, 1999). Card et al. argue that visualizing an information space allows us to use our highly evolved visual systems and processes, such as visual search, to find patterns and search for entities in more abstract spaces, which they refer to as using vision to think. An example might be a researcher being easily able to find a trend or and outlier in data when they visualize their data as a 3-d scatter plot.

Another important contribution is the theory of Information Foraging (Pirolli and Card, 1999) which applies concepts from optimal foraging theory in animals to understand how humans search for information. This theory applies the ideas from optimal foraging theory (developed as a theory of how animals search for food) to how human users search for information.

Finally, psycholinguistic analyses of how people talk about searching information spaces reveal that they often use spatial metaphors. For example, they refer to web pages as places, and of going to their favorite web sites, although the language used to search information spaces has changed somewhat over time (Matlock et al., 2014).

In summary, I have given several disparate examples of how cognitive psychologists have thought about aspects of "spatial search." I hope that the discussions at this specialist meeting will help us understand what is similar and different about these instances of human search and that insights from cognitive psychologists at this meeting (several of whom are cited here) can inform comptuational issues in search and how to design human-computer interfaces for search of information spaces.

References:

Card, S. K., Mackinlay, J., and Schneiderman, B. (1999) *Readings in information visualization: Using vision to think*. San Francisco: Morgan Kaufmann.

Cheng, K. (1986). A purely geometric module in the rat's spatial representation. *Cognition 23*(2): 149–178.

Cheng, K., and Newcombe, N. S. (2005). Is there a geometric module for spatial orientation? Squaring theory and evidence. *Psychonomic bulletin and review 12*: 1–23.

Eckstein, M. P., Drescher, B., Shimozaki, S. S. (2006) Attentional cues in real scenes, saccadic targeting and Bayesian priors. *Psychological Science 17*: 973–980.

Godden, D. R., and Baddeley, A. D. (1975). Context-dependent memory in two natural environments: On land and underwater. *British Journal of psychology 66*(3): 325–331.

Matlock, T. Castro, S. C, Fleming, M., Gann, T. M. and Maglio, P. (2014). Spatial metaphors of web use. *Spatial Cognition and Computation 14*: 306–320.

Pirolli, P., and Card, S. (1999). Information foraging. *Psychological review 106*(4): 643

Todd, P.M., and Gigerenzer, G. (2007). Environments that make us smart: Ecological rationality. *Current Directions in Psychological Science* 16(3): 167–171.

Geographic Information Spaces

STEPHEN C. HIRTLE
School of Information Sciences
University of Pittsburgh
Email: hirtle@pitt.edu

The relationship between physical space and human cognition has a long history of relating the impact of spatial cues on information retrieval in multiple ways. In this position paper, three distinct problems are considered in terms of the how the mind organizes and retrieves information by location. Each of these is discussed, in turn, below.

First and foremost, spatial location is one feature that can be automatically encoded with other information, often, but not always, without effort [5]. This is hardly a new finding, as spatial location has long been known as a foremost organizer for memory going back at least to the Method of Loci. This method described by the ancient Greeks demonstrated that a list of items could be most easily memorized by imagining the items in specific locations along a path. In a more modern version, one often finds themselves able to recall the specific location where a specific section of an audio book was heard while driving. This kind of memory can also been seen in recent studies showing that there may be memory for the computer folder of containing information, but not the information itself [8].

Thus, there is strong evidence to suggest the encoding of spatial location of information is both automatic and enduring. These effects are often amplified by when one returns to the original location where the information was first encoded, as both verbal and non-verbal cues can result in the recall of what appeared to be long-lost information. Finally, these ideas have been integrated into various memory tools, such as location-based reminding systems [4].

In a very different domain, there has been much research on the use of spatial ontologies for geographic information retrieval, when we want to generalize spatial or conceptual terms to improve the search result [3]. For example, a query for "rivers" that can be used as transportation routes might be expanded semantically to include canals, but not lakes and ponds. Alternatively, a query for amusement parks in Pittsburgh might be expanded geographically to include Kennywood, a historic American amusement park that is in the Pittsburgh area, but not strictly within the Pittsburgh city limits. The problem of which dimensions (semantic or spatial) to generalize upon is often quite difficult, as it is often the concept itself or the motivation of the user that determines how such terms should be generalized. Janowicz, Raubal, and Kuhn [3] suggest including a context that restricts the domain of application to those types that share the functional feature of interest. Similarity is then defined for within the restricted feature set.

Finally, there is a set of interesting methodological questions. Exactly how to compute the similarity is a long-standing problem of the clustering and scaling literature. Almost three decades ago, Pruzansky, Tversky, and Carroll [6] demonstrated how object similarity might be best modeled using a spatial representation, such as multi-dimensional scaling, in some cases and hierarchical tree, such as additive clustering, in others. The best method depended on whether the underlying features were continuous or discrete, which is reflected in the numerical judgments of similarity. The implication is that spatial representations should not be used in an arbitrary manner and that some attempts of spatialization can be misleading without taking care not to overgeneralize the meaning of the derived dimensions [1].

Together, the message is that memory and space are interconnected and that the memory for geographical notions may best be represented in a space that reflects the structure of the knowledge [2, 7]. This space of geographic information is complementary to geographic space that holds the information. Thus, the term "Geographic Information Space" may more accurately capture this notion, despite the seeming redundancy in the term.

[1] Hirtle, S. C. (1995). Representational structures for cognitive space: Trees, ordered trees and semi-lattices. In *Spatial Information Theory a Theoretical Basis for GIS*, pp. 327–340. Springer Berlin Heidelberg.
[2] Hirtle, S. C. (2011). Geographical design: Spatial cognition and geographical information science. *Synthesis Lectures on Human-Centered Informatics 4*(1): 1–67.
[3] Janowicz, K., Raubal, M., & Kuhn, W. (2011). The semantics of similarity in geographic information retrieval. *Journal of Spatial Information Science* 2: 29–57.
[4] Ludford, P. J., Frankowski, D., Reily, K., Wilms, K., & Terveen, L. (2006, April). Because I carry my cell phone anyway: functional location-based reminder applications. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, pp. 889–898. ACM.
[5] Naveh-Benjamin, M. (1987). Coding of spatial location information: An automatic process? *Journal of Experimental Psychology: Learning, Memory, and Cognition 13*(4): 595.
[6] Pruzansky, S., Tversky, A., & Carroll, J. D. (1982). Spatial versus tree representations of proximity data. *Psychometrika*, *47*(1): 3–24.
[7] Skupin A. and Fabrikant, S.I. (2007) Spatialization. In: Wilson, J. and Fotheringham, S. (Eds.). *Handbook of Geographic Information Science.* Blackwell Publishers.
[8] Sparrow, B., Liu, J., and Wegner, D. M. (2011). Google effects on memory: Cognitive consequences of having information at our fingertips. *Science 333*: 776–778.

Spatial Search

JON JABLONSKI
Department Head, Map and Imagery Laboratory
Davidson Library
University of California, Santa Barbara
Email: jonjab@library.ucsb.edu

Libraries maintain a large number of computational search tools for their users. From the end-users' point of view, these are mostly discovery engines—software products designed to help the user find the information they are seeking. Today, because most information is consumed digitally, these same systems offer access to the users as well. Much like Google, users type words into a box, press search, and are given result sets based on complex algorithms that are largely opaque to the users.

While it may be true that there is a spatial component to most information, in operational bibliographic systems, true geospatial search is mostly absent.

For my context, true geospatial information search:

- takes into account an extent
- requires a system that understands the spatial hierarchies of toponymy

Either one of these technically challenging, but doing both at the same time is proving to be elusive. Why is this?

Spatial search that takes extent into account is in widespread use in geoportals--all of the various search engines that provide access to maps and spatial data. The end-user manipulates an interactive web map to an area of interest, enters a keyword or two, and search results are ranked accordingly. Spatial extent is simultaneously an easy to understand concept for the user, and a metadata element that drives search results. This search paradigm is seen mostly in tools designed to store and provide access to spatial information. OpenGeoPortal in academic libraries, data.gov at the US federal level, and the European Union's INSPIRE geoportal all use this paradigm.

In bibliographic databases (search engines that provide discovery of books, journal articles, etc.), spatial search is often connected to specific controlled vocabularies. Mainstream library catalogs typically lump spatial vocabularies together with other subject-oriented vocabularies. American universities almost universally use the Library of Congress Subject Headings, which has a rich vocabulary of place names that are combined with other terms, such as the subject phrases:

Environmental impact analysis—Illinois—Peoria
Housing—Illinois—Peoria Metropolitan Area—Statistics
Public buildings—Illinois—Peoria County

Unfortunately, systems index these terms simply as text, ignoring all of the semantic meaning attached to the words. From the point of view of the system, Peoria is an equivalent concept to Statistics, and the relationship between the three Peorias is obscured except for the most skilled of system users. Separating the concepts semantically is the job of the end-user.

In some arenas, the situation is actually degrading. Even though the American Geosciences Institute's GeoRef database contains a spatial subject field, my library's vendor of the database, ProQuest, has combined place name subject terms with topical terms--resulting in Holocene and Peoria receiving equal meaning and weight. Systems that segregate place names into a discretely searchable field are actually disappearing.

Fortunately, the world is organized in such a way that *most* of the time, for *most* searches, Google-style free-text searching *mostly* returns satisfactory results. Librarians call this satisficing, but many of us think we can do better. Moreover, as the volume of digital data increases, and the variety of formats that libraries manage explodes, users are more frequently encountering the outer limits of what systems can return. At the same time, the users' expectations continue to rise. Encouraged by the location based services available on their mobile phones, the widespread use of "slippy maps," and media that portrays digital search technologies as magical, users expect discovery systems to perform better.

The limitations inherent in this ecosystem of bibliographic tools are very well described. A plethora of research projects have provided computational and knowledge organization techniques to apply to this problem domain. For example, for at least fifteen years, semantic web researchers have been developing linked data standards and tools, and a huge community is emerging that marks up texts and data—both manually and through automated techniques.

However, the tools most commonly used in academic libraries continue to lag behind. Why?

Careful analysis from a socio-technical standpoint can help to answer this question. Qualitative research methods from sociology and anthropology, as manifested in a field often dubbed Science and Technology Studies (STS), are increasingly being applied to help solve problems in technology design. STS maps the intersection of technology and culture. It treats the various factors of system design as independent actors in a network (hence actor-network theory). One strain of thought in STS, often associated with Bruno Latour, treats human and non-human actors equally: hardware and software artifacts, datasets, laboratories, and universities are treated just the same as programmers, system administrators, scientists, and librarians. If there are actors preventing the uptake of better spatial search technologies in library discovery services, actor-network analysis should be able to identify them.

The research question implied by this argument is not simply "What are the socio-technical impediments to spatial search in academic libraries?" Rather, it is: "Given the complex socio-technical factors at play, how can academic libraries and the research enterprise work together to enhance each other's work?"

I plan to participate in this Specialist Meeting in order to closely observe how spatial search research questions are developed, and to present the perspective of the operational end of the university. As research projects develop, how can the library serve as both resource for, and beneficiary of, the research? How can we best encourage research that results in tools that we can apply toward operational goals?

Semantic Signatures for Places of Interest

KRZYSZTOF JANOWICZ
STKO Lab
Department of Geography
University of California, Santa Barbara
Email: janowicz@ucsb.edu

BENJAMIN ADAMS
Centre for eResearch
Department of Computer Science
University of Auckland, New Zealand
Email: 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 park. 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 hotel 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 go to any bar, they favor certain types of bars that cannot be distinguished just by properties such as Wi-Fi 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., and Card, S. (1999). Information foraging. *Psychological Review*, 106(4): 643–675.

[2] Backstrom, L., Kleinberg, J., Kumar, R., and 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., and 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., and 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., and 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.

What Makes Things Searchable is the Underlying Scaling

BIN JIANG
Professor, GeoInformatics and Computational Geography
University of Gävle, Sweden
Email: bin.jiang@hig.se

We search on a daily basis; for example, to find books in a library or to find information on the Internet. The former occurs in well-organized databases, while the latter—commonly conducted through a search engine like Google—takes place via self-organized or unstructured databases. This difference can be seen from the statement "*We don't search for information, we 'Google' it*" (Auletta 2009), which implies a paradigm shift in search. In line with the Google-like search, the present paper attempts to argue that what makes things searchable is the underlying scaling. In other words it is the underlying scaling property of spatial or non-spatial things that makes things searchable.

Scaling refers to the pattern or hierarchy that underlies the notion that there are far more small things than large ones, or equally, far more unpopular things than popular ones, and far more meaningless things than meaningful ones. We rank things according to their geometric, topological and/or semantic property in decreasing order and divide all the things around an average into two unbalanced parts: a minority of large things above the average in the head, and a majority of small things below the average in the tail with respect to the rank-size plot (Zipf 1949). This dividing process continues recursively for the large things in the head until the notion of far more small things than large ones is violated. This recursive process, guided by head/tail breaks (Jiang 2013a), ends up with different classes or hierarchical levels for data with a heavy-tailed distribution, such as power laws, lognormal, and exponential distributions. In terms of searches, the fact that there are far more unrelated things than related ones means that something like the Internet, or the World Wide Web, is searchable. For example, the first page of a Google search results page shows the most related web pages, followed by uncountable less related pages. The most related things are unique and can be distinguished from the abundant others. Shown in the rank-size plot, the most related things are in the head, while the abundant others are in the tail.

Things are searchable as long as they bear the scaling property. A human face is searchable because it contains far more small things than large ones; notably, the eyes and mouth receive far more attention than anything else (Yarbus 1967, Figure 1). A city is searchable because of far more small city artifacts than large ones, or far more low-density areas than high-density ones (Salingaros 2005, Jiang 2013b). In this regard, there is not much difference between a city and a human face in terms of forming their mental maps (Haken and Portugali 2003). Scaling is also the first and foremost reason why an image of a city or mental map can be easily formed in human minds (Jiang 2013b). A set of interconnected cities is searchable because there are far more small cities than large ones, which is an example of Zipf's law (Zipf 1949). Both a city and a set of interconnected cities are self-organized, but they are searchable because of the underlying scaling. In general terms, geographic space is searchable, because it bears the scaling property. Drawing on our previous work on natural cities (Jiang and Miao 2014), I will use massive number of tweets as an example with which to elaborate how to uncover the underlying scaling property, or make tweets searchable. My elaboration will surround the three suggested perspectives: geospatial, computational, and cognitive.

The geospatial perspective is what underlies the notion of spatial search to differentiate it from non-spatial search. The term geospatial refers not only to spatial, but also temporal elements in the process of information search. When analyzing several million tweets, their locations and time stamps are critical for spatial search and for uncovering the underlying patterns. The concept of natural cities refers to human settlements aggregated naturally and objectively from massive tweets locations based on head/tail division rule – the non-recursive version of head/tail breaks (Figure 1). Because there are far more tweets in high-density locations than in low-density locations, the high-density locations that are above the average can be aggregated to form individual natural cities. The natural cities again bear the scaling property: far more small natural cities than large ones. The scaling hierarchy of natural cities can be an important index for hierarchically organizing tweets, which will significantly benefit spatial search.

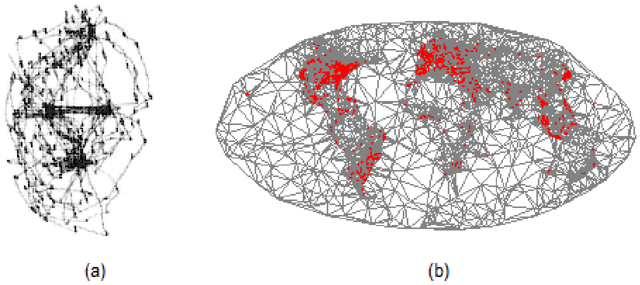

Figure 1: The searchable elements in the human face and earth surface
(a) fixations from eye-watching the human face (Yarbus 1967), and
(b) natural cities (red spots) derived from 12-hour tweets locations

The generated natural cities, which are bottom-up in nature, provide a better alternative than conventional cities or any top-down imposed geographic boundaries. The natural cities are delineated using a unique cutoff (or the average) for the entire world, which makes them more objective or natural than the conventional cities or administrative boundaries. Secondly, the time stamps enable us to study the evolution of natural cities (Jiang and Miao 2014), or that of real cities defined on populations, at a more fine temporal scale than any conventional counterparts. Because of these two advantages, I believe that natural cities are better than conventional cities for spatial indexing and for designing search functionality in large information systems. The natural cities are derived from all tweets locations, and, in turn, they can serve back to all the tweets for spatial search. Computationally, the process of deriving natural cities is very intensive, because it involves all tweets locations. Given the scaling hierarchy, geographic representation should adopt a topological representation like the web graph behind Google (Page and Brin 1998). Unfortunately, current geographic representations, which are essentially geometric in terms of geometric locations, distances, and directions, prevent us from seeing the underlying scaling. The topological representation would help considerably for spatial indexing and spatial search.

The concept of natural cities has some implications for cognitive mapping. No one can deny that the fixations of the eye movement captured from watching the photographed face are not about a human face (Figure 1). The fixations as a whole capture some essential elements or landmarks of a human face. Equally, no one can deny that the map of natural cities is not a map of the world. In fact, one can simply zoom into individual natural cities to find out that the natural cities quite accurately capture their corresponding real cities. Although it is debatable, I believe that the computation of the natural cities, from the unique geospatial perspective, resembles that of cognitive mapping, which is essentially based on the head/tail breaks process. Search or spatial search in particular deserves further research. Accordingly, I look forward to attending the specialist meeting to engage in the discussions.

References:

Auletta K. (2009), *Googled: The End of the World as We Know it*, The Penguin Press: London.
Haken H. and Portugali J. (2003), The face of the city is its information, *Journal of Environmental Psychology* 23: 385–408.
Jiang B. (2013a), Head/tail breaks: A new classification scheme for data with a heavy-tailed distribution, *The Professional Geographer* 65(3): 482–494.
Jiang B. (2013b), The image of the city out of the underlying scaling of city artifacts or locations, *Annals of the Association of American Geographers* 103(6): 1552–1566.
Jiang B. and Miao Y. (2014, accepted), The evolution of natural cities from the perspective of location-based social media, *The Professional Geographer* xx(xx), xx–xx. Preprint: http://arxiv.org/abs/1401.6756
Page L. and Brin S. (1998), The anatomy of a large-scale hypertextual Web search engine, *Proceedings of the Seventh International Conference on World Wide Web*, 107–117.
Salingaros N. A. (2005), *Principles of Urban Structure*, Techne: Delft.
Yarbus A. L. (1967), *Eye Movements and Vision*, Plenum Press: New York.
Zipf G. K. (1949), *Human Behaviour and the Principles of Least Effort*, Addison Wesley: Cambridge, MA.

Toward High Resolution Spatial Search: From Documents to Spatial Facts

CHRISTOPHER JONES
Professor, Geographical Information Systems
School of Computer Science and Informatics,
Cardiff University
Email: c.b.jones@cs.cf.ac.uk

This paper highlights limitations of current geographical web search methods for indexing and accessing content in natural language documents and discusses research directions for automated extraction and indexing of geo-coded spatial relations for exploitation on the semantic web.

Geographical information resources on the web take a variety of forms ranging between highly structured, partially structured and relatively unstructured. Structured resources include the digital maps and associated databases of conventional GIS, accessible via web services, and the semantic web structured resources such as the geographical components of DBpedia and Freebase, that can be queried with GeoSPARQL. Semi-structured resources include social media that typically combine structured data objects with free text reviews that can be accessed with dedicated APIs. Unstructured resources are typified by web documents consisting of text and other media within HTML documents. The latter unstructured resources contain a vast wealth of geographical information embedded in natural language texts that remain only crudely searchable using current commercial search engine methods. These methods are subject to the hit and miss of keyword search in which the user is presented with a list of documents that contain the search terms. Some of the documents may be relevant but the user is still expected to browse through them and read the respective content to determine whether that is the case. Focused answers to geographical web search, provided for example in response to some queries on Google, depend largely on access to structured resources most of which have been created with some manual input, as in the case of GIS or in creating tables in Wikipedia that have been transformed automatically to semantic web resources (DBpedia). These semantic web resources do facilitate direct answers to queries, but are limited in their scope relative to the amount of knowledge represented in other forms, as well as being subject to limitations of geo-spatial referencing, consistency and vagueness. The challenge remains to develop effective methods to extract knowledge from natural language resources and to find methods to index and access it in an effective manner.

In recent years there has been a proliferation of research publications on spatial keyword search and spatio-textual indexing (SK/ST) methods, dedicated to assisting in the process of searching geographic content on the web (e.g., Wu et al). These methods combine inverted text indexing and spatial indexing with various degrees of integration. The spatial indexing components depend upon geo-parsing and geo-coding techniques that detect the presence of toponyms within text documents and relate the toponyms to a quantitative coordinate-based geo-spatial reference. This approach to spatio-textual indexing was introduced more than 10 years ago and the following research developments have been largely confined to improving their computational efficiency, especially with regard to the application of associated spatial and textual relevance ranking procedures. From the user's point of view however they provide minimal improvement in the quality of the information returned. This is due to the fact that documents are treated as a bag of words, and hence while the methods can retrieve a list of documents that contain the keywords that relate to the spatial query footprint, they provide only limited assurance that there is any logical connection between them, as the query terms (keywords) could be employed quite independently of each other within an individual document. Furthermore the methods cannot by themselves process queries that require direct answers to specific geographic questions or geospatial aggregation of information, as they only retrieve complete documents.

Evaluations of performance of these types of systems such as in the GeoCLEF competitions have found that the SK/ST indexing methods provide relatively little advantage over standard search engine methods (based on inverted files in combination with a mix of relevance ranking methods). Other campaigns such as the MediaEval Placing Task have helped to progress the development of novel methods to categorise the geographical scope of web documents, which while it helps to increase the number of documents that can be indexed with ST/SK indexing methods, it again provides little progress in improving the quality of the user experience for geo-spatial search.

In attempting to index documents with geographical information in a manner that provides increased assurance of the logical association between information and its geographic context, it is possible to identify different levels of granularity at which this may be accomplished. At the finest level, natural language processing (NLP) methods can be used to detect relations between spatial and non-spatial entities, between spatial entities, and between non-spatial entities, which can then be encoded as factual statements, with a subject, predicate and an object. Such triples could be indexed directly via for example GeoSPARQL triple stores or other forms of spatially enabled databases. It is also possible to envisage expanding their source natural language document context to topically coherent paragraphs that had a defined geographic context, and hence could be treated as document chunks for indexing purposes in a spatio-textual inverted index. Progress has been made with spatial relation extraction techniques (Kordjamshidi et al.) that associate concept terms with spatial objects via a spatial predicate, but such work is still at an early stage in that the spatial context of the relations still needs to be grounded (geo-referenced) with regard both to the location of spatial objects and the interpretation of vague spatial prepositions that may serve as the predicates (Bateman et al).

While considerable progress has been made in toponym resolution in documents, there are still challenges in increasing the granularity at which spatial objects are detected and grounded. This relates particularly to the need to access spatially-referenced information about buildings, to support applications relating to cultural heritage, management of industrial plant, and support for emergency services and the visually impaired. In the cultural heritage domain alone there are very many detailed tourist guides to historic buildings, architectural texts, and historic topographic narratives describing the spatially-specific details of buildings and their component features (doors, windows, towers, sculptures etc.) and associated events. Here named objects may need to be disambiguated with respect to their parent building, while spatial grounding requires determining the frame of reference of localization expressions (Tenbrink), as well as their actual spatial applicability, e.g. with a density field (Hall et al). For some applications the spatial referencing needs to be in 3D. There are an increasing number of 3D models of buildings in for example Trimble 3D Warehouse, most of which lack any semantic annotation. However their geometry in combination with texture-mapped images, as well as point cloud 3D models derived from multiple photos, provide the potential basis for spatial grounding of facts about the buildings that have been extracted from other sources (Russell et al.).

Assuming that spatial relations can be extracted and represented in a form analogous to that of existing geo-spatial semantic web content, in which objects are explicitly referenced to geographical coordinates, major challenges remain with regard to searching this structured knowledge. These concern for example coarse granularity geo-referencing (with spatially extensive objects often referenced to point locations), interpretation of vague spatial relations that may refer to an explicitly geo-coded object, inconsistency of multiple conflicting representations of the same real world phenomena and the integration of qualitative and quantitative descriptions of these phenomena.

References

Bateman, John, Hois, Joana, Ross, Robert J., Tenbrink, Thora, 2010. A linguistic ontology of space for natural language processing. *Artificial Intelligence* 174: 1027–1071.

Hall, M.M., P.D. Smart and C.B. Jones, 2011. Interpreting spatial language in image captions. *Cognitive Processing* 12(1), pp 67–94.

Kordjamshidi, P., M. van Otterlo, and M.F. Moens, 2011. Spatial role labeling: Towards extraction of spatial relations from natural language. ACM Transactions on Speech and Language Processing, 8(3).

Russell, B. C. et al., 2013. 3D Wikipedia: Using online text to automatically label and navigate reconstructed geometry. ACM Transactions on Graphics*,* 32(6).

Tenbrink, T., 2011. Reference frames of space and time in language. *Journal of Pragmatics 43* (2011): 704–722.

Wu, D., G. Cong, and C.S. Jensen, 2012. A framework for efficient spatial web object retrieval. *The VLDB Journal* 21(6): 797–822.

Place Reference in Text as a Radial Category:
A Challenge to Spatial Search, Retrieval, and Geographical Information Extraction from Documents that Contain References to Places

ALAN M. MACEACHREN
Professor,Geography
Affiliate Professor, Information Sciences & Technology
Director, GeoVISTA Center
The Pennsylvania State University
Email: maceachren@psu.edu

The concept of place is fundamental to spatial search as it relates to finding geographical information in heterogeneous information repositories as well as to the kinds of "spatial search" that are central to human wayfinding in the world. This position statement focuses on place in the context of spatial search for and within text documents, but the conceptualization of place proposed has implications for spatial search more broadly. Spatial search in text documents requires solving a range of challenges related to place, including: recognizing place entity mentions in text, disambiguating and locating the place entities, and determining whether the document is "about" the place(s) mentioned or those place mentions serve some other purpose.

Research over the past several years has provided a base from which to address each of these questions (Zhang et al. 2009; Zhang et al. 2010; Jaiswal et al. 2011; Zhang et al. 2012; Karimzadeh et al. 2013; Xu et al. 2014; Wallgrün et al. submitted). Based upon that research, I propose that one of the most fundamental issues in spatial search focused on retrieving and extracting geographically-relevant information from documents is understanding and being able to recognize what constitutes a reference to a place. My argument here is that place reference is best conceptualized as a *radial category*, but that current entity recognition approaches take a classical category theory approach to recognizing place references. This existing approach impedes strategies for spatial search in text as well as those for geoparsing and mapping the content of documents.

It is important to distinguish between "place" and "place reference." Conceptualization of place is also a fundamental research question (Goodchild 2011) and category theory (within which the concept of radial categories is defined) has been applied successfully to investigate conceptualization of place (Mark, Smith, and Tversky 1999; Lloyd, Patton, and Cammack 1996). Here, however, my focus is on linguistic or textual references to place and how to recognize and interpret them. Thus, the focus is not on what a place is conceptually or in practice (although conceptualization of place is important in understanding place references) but on distinguishing references to place from references to people, organizations, and other entities.

Radial categories (Lakoff 1987), are structured with respect to a prototype and include members in relation to some "center" or prototype (Janda 2010). While I am not aware of any empirical research focused on identifying prototype place references (there is work on basic level geographic features that includes place concepts, [e.g., Lloyd, Patton, and Cammack 1996]), nor the characteristics of "place reference" as a category, I can speculate that London or Paris might constitute an exemplar, thus close to a prototype place reference for many people. But, someone from China might center their conceptualization of place on Beijing while someone from India might have New Delhi as their prototype. On the other hand, an Amish farmer from the eastern U.S. might center their concept of place on the home farm or local livestock market. Thus, the category of place differs for individuals from different places and who have a different pattern of behavior in the world. Whatever the prototype, radial categories include entities that are more or less representative, thus more or less central or peripheral to the prototype. The position in this radial category space has implications for strategies used to recognize place references in text and for how such references are tagged for subsequent use.

For the category of "place references" (in text or speech), centrality (and thus the likelihood that a speaker/writer, in a statement containing a place name or related reference, is specifically thinking of an entity as a "place" versus another entity type) is likely to depend upon a range of factors. As part of recent research to develop tools for place reference corpus building and then to apply the tools to building a Twitter place reference corpus (Wallgrün et al. submitted), several factors were identified as responsible for disagreements among human annotators (given the task of tagging place references in tweets). Drawing on this work and further analysis of our tweet-tagging results, I propose the following as factors that influence whether a named entity was intended as a place reference; the combination of factors will determine relative position of a reference within or outside the radial place reference category:

Geographical scale: Research has demonstrated that bonds people form with places differ according to geographical scale (e.g., neighborhood, city, country) (Bernardo and Palma-Oliveira 2013). The intensity of place identity and attachment is likely to influence the extent to which a named entity is conceptualized as a "place."

Use as a noun versus adjective: Names as nouns are more certainly place references than names used as an adjective. But, use as a possessive adjective is intermediate. Given three endings of the phrase "I'm going to {the Pirates game in Pittsburgh; Pittsburgh's Pirates game; the Pittsburgh Pirates game}", the first is clearly a place reference and the last is clearly a reference to a sports team (that is from a place). The middle case, however, can be interpreted as a reference to the place, which has a sports team for which only the nickname is given.

Precision of reference: The distinction between an organization and an instance of that organization that is a place is a function of precision of reference. The statement "I often drink Starbucks coffee" refers to the company while "We usually stop at Starbucks" might be a reference to a particular coffee shop or to the chain generally. In contrast, "let's meet at the Starbucks in the hotel lobby" is more certainly a reference to a (within building) place.

Interdependence: Entities can be clear place references in one context and a qualifier of a separate place reference in another. For example, in the two version of this sentence: "I'm traveling to {Illinois; Springfield, Illinois}, the first refers to the state as a distinct place while the second uses "Illinois" to clarify which Springfield is intended (with Illinois secondary).

Agency or lack of it: Place names can be used in ways that imply agency; in these cases, there is a tension between conceptualizing the reference as a place versus as an agent exhibiting behavior. An example is: "Canada Takes a Wait-And-See Approach to New Cage Regulations," Alberta Farmer, Posted Nov. 7, 2011 by Sheri Monk in Livestock; accessed Sept. 19, 2014 (http://www.albertafarmexpress.ca/2011/11/07/canada-takes-a-waitandsee-approach-to-new-cage-regulations/).

This statement only scratches the surface of the potential application of cognitive category theory to conceptualization of place in the context of spatial search. There is a range of past research on application of category theory to understanding how geographical features are conceptualized (e.g., Mark, Smith, and Tversky 1999; Usery 1993) and newer work on application to contexts such as tagging Flickr photos (e.g., Stvilia and Jörgensen 2010). A next step is to focus on integrating knowledge about how place and related geographic features are conceptualized with that on how place is referred to in text. This integration has the potential to enhance our ability to search for documents that are "about" places of interest and to recognize, extract, disambiguate, and locate the references. One objective is to develop entity recognition methods that take context into account and that can recognize different kinds and intensities of place reference so that spatial search for geographic information in text documents can be better tailored to particular use contexts.

References:

Bernardo, F., and J. Palma-Oliveira. 2013. Place identity, place attachment and the scale of place: The impact of place salience. *Psyecology* 4(2): 167–193.

Goodchild, M. F. 2011. Formalizing Place in Geographic Information Systems Communities, Neighborhoods, and Health. In *Communities, Neighborhoods, and Health*, eds. L. M. M. Burton, S. A. P. Matthews, M. Leung, S. P. A. Kemp and D. T. T. Takeuchi, 21–33: Springer New York.

Jaiswal, A., X. Zhang, P. Mitra, S. Pezanowski, I. Turton, S. Xu, A. Klippel, and A. M. MacEachren. 2011. GeoCAM: A Geovisual Analytics Workspace to Contextualize and Interpret Statements about Movement. *Journal of Spatial Information Science* 3: 65–101.

Janda, L. A. 2010. Cognitive linguistics in the year 2010. *International Journal of Cognitive Linguistics* 1 (1): 1–30.

Karimzadeh, M., W. Huang, S. Banerjee, J. O. Wallgrün, F. Hardisty, S. Pezanowski, P. Mitra, and A. M. MacEachren. 2013. GeoTxt: A Web API to Leverage Place References in Text In *7th ACM SIGSPATIAL Workshop on Geographic Information Retrieval*. Orlando, FL: ACM.

Lakoff, G. 1987. *Woman, Fire, and Dangerous Things: What Categories Reveal about the Mind*. Chicago: University of Chicago Press.

Lloyd, R., D. Patton, and R. Cammack. 1996. Basic-level geographic categories. *Professional Geographer* 48(2): 181–194.

Mark, D. M., B. Smith, and B. Tversky. 1999. Ontology and Geographic Objects: an empirical study of cognitive categorization. Paper read at COSIT'99—Conference on Spatial Information Theory, at Berlin.

Stvilia, B., and C. Jörgensen. 2010. Member activities and quality of tags in a collection of historical photographs in Flickr. *Journal of the American Society for Information Science and Technology* 61(12): 2477–2489.

Usery, L. 1993. Category theory and the structure of features in geographic information systems. *Cartography and Geographic Information Systems* 20(1): 5–12.

Wallgrün, J. O., M. Karimzadeh, A. M. MacEachren, F. Hardisty, S. Pezanowski, and Y. Ju. submitted. Construction and First Analysis of a Corpus for the Evaluation and Training of Microblog/Twitter Geoparsers.

Xu, S., A. Klippel, A. M. MacEachren, and P. Mitra. 2014. Exploring regional variation in spatial language using spatially-stratified web-sampled route direction documents. *Spatial Cognition and Computation*.

Zhang, X., P. Mitra, A. Jaiswal, S. Xu, A. M. MacEachren, and A. Klippel. 2009. Extracting route directions from webpages. Paper read at Twelfth International Workshop on the Web and Databases (WebDB 2009) June 28, 2009, at Providence, Rhode, Island, USA.

Zhang, X., P. Mitra, A. Klippel, and A. MacEachren. 2010. Automatic Extraction of Destinations, Origins and Route Parts in Human Generated Driving Directions. Paper read at GIScience 2010, at Zurich, Switzerland.

Zhang, X., B. Qiu, S. Xu, P. Mitra, A. Klippel, and A. M. MacEachren. 2012. Disambiguating Road Names in Text Route Descriptions using Exact-All-Hop Shortest Path Algorithm. Paper read at European Conference on Artificial Intelligence (ECAI), Aug. 27–31, at Montpellier, France.

Spatial Search

TEENIE MATLOCK
Founding Faculty, Cognitive and Information Sciences
School of Social Sciences, Humanities and Arts
University of California, Merced
Email: tmatlock@ucmerced.edu

A long, long line of research has investigated metaphor. For centuries, researchers in philosophy, linguistics, cognitive psychology, literature, anthropology, education, and other fields have discussed what metaphor is and how humans make sense of it. In recent years, there is consensus that metaphor is nothing magical or special. It is not in fact reserved only for poetry or literary work. Rather, it is part of all language and part of our basic conceptual system (Gibbs, 1994; Lakoff and Johnson, 1980).

Spatial metaphors are especially important to everyday thought and communication (Lakoff and Johnson, 1980). They figure into how people talk about all sorts of things the world, including time (Clark, 1973; Moore, 2014), social relationships (Matthews and Matlock, 2011), romantic love (Gibbs and Nascimiento, 1996), emotions (Kövecses, 2000), politics (Lakoff, 1997; Matlock, 2012), physics (Pulaczewska, 2011), numbers and mathematics (Núñez, 2011), health and illness, especially cancer (Reisfield and Wilson, 2004), sustainability and environmental issues (Larson, 2011), and technology, including the Web and other information spaces (Maglio and Matlock, 1999).

In all these cases and more, people use their knowledge of and experience with physical space (the source domain), including motion, to structure how they talk and think about something relatively abstract (the target domain), i.e., unfamiliar, complex, obscure, or intangible. (Imagine how hard it is to talk about emotions or the web without spatial language and thought!) For instance, when early personal computer users first attempted to search for information on the World Wide Web, web, they naturally used spatial metaphors. This was seen, and still is today, in statements like, "I started to go to your website but decided to come back to *Yahoo!*" In such language, motion verbs (*go, come*) are used to describe information search, not actual physical motion. Nobody is actually going anywhere (Maglio and Matlock, 1999; Matlock, Castro, Fleming, Gann, and Maglio, 2014). Similarly, in talking about math, people use spatial metaphors, such as the number line, in which numbers are viewed on a left to right horizontal array, or they use a vertical schema, in which numbers are thought of as higher or lower relative to each other (e.g., "9 is higher than 7") (see Lakoff and Núñez, 2000). And in discussing politics, people use spatial language. For instance, in talking about political candidates and their status relative to each other right before an election (Matlock, 2012; 2013), people use motion metaphors (e.g., "Obama sprinter ahead of Romney," "The candidates are inching toward the finish line").

Much of the research on metaphorical thought focuses on how metaphors are comprehended. Often questions have often focused on the extent to whether and how metaphors are derived from literal language, and whether metaphorical language, or more generally, figurative language, is more difficult to interpret than literal (see Gibbs, 1994, for overview). Despite many claims about how spatial metaphors in particular are important to everyday thought, little research has been devoted to the dynamics of metaphorical thought or their utility in everyday reasoning. Is the mapping from a spatial source domain to some target domain always unidirectional, as some of the early metaphor research would have us believe? And if it is largely or entirely unidirectional, how does it work? What spatial information is carried over, and what is not? Also, what features do all spatial metaphors seem to share? For spatial metaphors that have motion in the source domain, for instance, TIME IS MOTION, is there mental simulation of motion? And last, when does space serve as the source domain in spatial reasoning, for instance, in reasoning about complex spatial layouts?

To get to a better understanding of how spatial metaphors work, we need a better handle on what spatial metaphors are and how they are actually used. We also need to investigate when and how metaphors are learned, how they naturally arise, and how they can be applied to help with real world problems, including environmental issues and political negotiations. This will require going beyond the study of metaphor comprehension in simple lab experiments. It will necessitate looking closely at how metaphors are used in natural discourse “in the wild”, across different contexts, across different slices of time, and across different cultures. It will also require looking at how linguistic information or interacts with or affects metaphor use, for instance, looking at how tense and aspectual information (e.g., “was “VERB+ed” versus “VERB+ed”) interacts with metaphor. And it will require *carefully* studying how and when metaphorical language is used in learning and discussing abstract material.

Expanding WorldMap

STEPHEN MCDONALD
Center for Geographic Analysis
Harvard University
Email: stephenmcdonald@cga.harvard.edu

with
BEN LEWIS
Senior GIS Specialist, Center for Geographic Analysis
Harvard University
Email: blewis@cga.harvard.edu

Finding spatial information using search engines is still very difficult. For this paper we focus on map services discovery, though there are many other kinds of spatial information on the web and elsewhere that need to be more discoverable.

The Harvard's Center for Geographic Analysis (CGA) is engaged in an NEH-funded effort to address the problem of map service discovery by extending the information gathering and search capabilities of an existing map publishing platform, WorldMap. The goal is to build a registry of all known map services and make it available it available via a public query API. WorldMap users will use this API to find layers in WorldMap and to find layers outside WorldMap. External layers that are used within WorldMap will accrue usage statistics which will be used to improve search. Uptime statistics will be gathered on all remote services to attempt to address the problem of the ephemerality of map services. Within WorldMap, feature level information will be used to improve search. A faceted interface will be provided which will include a layer footprint heat map of all search results even when results are in the millions. Any map client outside WorldMap will have access to the API to use to find and bind to all available map services.

Spatial and textual search on WorldMap and external services will be enhanced. Layer service metadata, such as it is, will be used to create a Solr index holding title and abstract, geodetic bounding box and feature level names and attributes. We also hope to use feature-level geometry to compute data density. WorldMap search algorithms we be based on this author's previous work implementing search on OpenGeoPortal.

Before the web's resources can be made searchable (spatially or otherwise) they must first be discovered. We will accomplish this is by crawling the web. As part of my Masters thesis I explored crawling for both spatial and semi-spatial data. A web-wide crawl is straightforward using data from CommonCrawl.org. Initial explorations succeeded in identifying large numbers of HTML links to kml and kmz files. After indexing, these files could be searched and previewed in OpenGeoPortal. As part of enhancing WorldMap, we will crawl over 4 billion pages, searching for all the world's OGC and ESRI REST Image Service servers as well as other significant spatial data resources. Preliminary research suggests millions of map service layers are publicly available. We build on GeoNetwork's recently expanded harvesting infrastructure and support user submission of URLs. Individual sites and services holding spatial data can be periodically re-crawled to obtain new data and to gather uptime and reliability statistics.

WorldMap may collaborate with the Lucene/Solr community to add support for a heat map query. This will be used in two ways. As WorldMap ingests larger and larger amounts of data it becomes difficult for users to gain an overview of its holdings. The search interface only makes individual layers and maps discoverable. By adding a heat map display based on the number of layers in each grid cell users will be able to see where and how the layer footprints are clustering. This will be useful to guide the researcher during the search process as well as for collections management. The ability to quickly determine data densities will also help us improve the ranking of search results. The ranking can depend in part on the number of features each layer has over the area of interest.

There is a problem of searchability that follows from enlarging the pool of discoverable maps and spatial datasets in a metadata weak environment: it is difficult to find material of interest in a list of thousands. Because volunteered data can be titled idiosyncratically and because contributors across disciplines do not share a common ontology for geospatial data, the growth in resources increases the difficulty of discovery. Even if all sites that made information available controlled the quality of its content, there would be no uniform quality standards. The problem is especially acute in an era of volunteered data. We thus plan to build the "wisdom of crowds" into the search mechanism so that over time search results will become more and more relevant as usage information grows and is used to weight the order of results returned. WorldMap users actively view, organize, rename, rank, comment on, and link to their data and the data of others. Over time usage statistics will improve search results not just for WorldMap users but for all users via the map service registry.

As part of thinking about better search UIs we are investigating client-side rendering of spatial data. Existing OGC protocols enabled web mapping and served the community well. However, with HTML5 and support for the canvas tag the client need not rely on a server to render spatial data into image tiles. Instead, the client can parse the raw data files and render them in JavaScript. As the map pans and zooms the client need not request new tiles, instead it can re-render the data as needed. On the browser, these graphic operations will leverage the host's GPU for higher performance. Client-based map drawing and data analysis code can be much more interactive since there is no latency from connecting to a server.

WorldMap is being used in many courses and for many ends. It currently has about 12,000 registered users, 14,000 data layers. About 1000 users/day from around the world use the system. WorldMap is in the process of being integrated with a social science archive called Dataverse to enable spatial visualization of its datasets. We are collaborating with HarvardX to support online course development which includes maps. WorldMap is also being integrated with Neatline to make it easier for Humanists to create fine grained narratives which include GIS layers.

Spatial Search

GRANT MCKENZIE
Doctoral Candidate, Department of Geography
University of California, Santa Barbara
Email: grant.mckenzie@geog.ucsb.edu

The term *Spatial Search* is loaded in that it caries with it multiple definitions, different meanings for various people across a range disciplines. The word S*earch* itself may reference physically searching, mentally searching or executing data searches online. The *Spatial* component of the phrase may speak to spatial objects in our natural environment, spatial relationships between entities or web documents tagged with specific geospatial locations (to name a few). In this position paper I take *Spatial Search* to mean the querying of (representations of) geospatial objects through an online service.

Spatial search has made significant advances in recent years from a number of different perspectives. Technologically speaking, spatial search has benefited from the ubiquity of location-enabled mobile devices driven by reduced cost and increased accuracy of technologies such as GPS. From a data perspective, the quantity of user-generated content has drastically enhanced existing gazetteers and given rise to new location-dependent online social networks (e.g., Foursquare). From a computational side, advances in spatial indexing, linked-data stores and geographic information retrieval have pushed the boundaries of what is possible with spatial search. Up until this point though, the majority of research has focused solely on the spatial or geospatial dimension of search ignoring additional and often highly influential dimension such as *time* and *theme*. Research should continue to pull on this thread of dimensionally-enhanced spatial search, asking how and to what degree concepts such as time and topic influence search and the way we interpret search results.

Take for example the task of "checking in," a concept made popular by the transportation industry, but more recently reflecting the act of publishing your platial location through an online social networking application. Computationally, *checking in* often involves taking a spatial location (set of geographic coordinates) obtained through a location-enable mobile device, and determining the platial location of the user controlling the device. In many ways this is an example of *spatial search* in its most basic form. Given a set of coordinates and a gazetteer of point of interests (POI), return a set of places ordered by Euclidean distance from closest to furthest away. Current *check-in* applications simply assume the user is at the closest place, returns the name of the venue, and the process is complete. While this works quite well in many cases, it makes the erroneous assumption that all places have the same probability of being visited at any time of day and day of the week. This methods implies that the only property of a place influencing an individual's decision to visit, is its spatial location.

In actuality, time plays a significant role in spatial search, and in particular this example of *checking in*. The probability of checking in to a *Night Club* at 8am on a Tuesday is much less than say a *Coffee Shop* and the opposite is true for Saturday at 1am [1]. There are specific temporal patterns and behaviors that are linked to human activities. The places that we chose to visit are often defined by the activities that the places afford. In many ways this is a logical extension of the socio-institutional affordances described by Raubal et al. [2]. The human construct of *Categories* has arisen from this need to group places by the allowed activities.

Leaving *Time* aside, additional dimensions such as *Theme* or *Topic* play an important role in both how spatial data is queried as well as how a resulting data set is interpreted. Let us explore another example specific to user-generated geo-content, but this time from a spatio-thematic standpoint. In many ways the value of *big data* lies in its variety. Large POI datasets generated through user contributed means (e.g., Yelp, Foursquare, Geonames) come with bias imposed by both the collection platform as well as the users contributing the data. By combining multiple datasets, much of the bias is reduced while increasing the breadth of content. For example, descriptive information for a specific POI such as price, noise level and wheel-chair accessibility may be present in the Yelp representation of the data, while number of check-ins, tips and rating may be contained in the Foursquare representation of the same POI. The results of a *Spatial Search* on a merged/conflated dataset would be of greater value than individual datasets alone.

Matching and conflating this data is an important, yet difficult task. In order to accomplish this task, dimensions other than the geospatial dimension must be explored. Simply matching user-generated POI between providers based solely on geographic coordinates has been shown to be less than adequate [3]. Inclusion of additional attribute dimensions is essential for increasing the accuracy of POI matching. By accessing thematic attributes of the data such as tips and reviews, POI can be compared in *topic* space based on the words, phrases and content contributed by individual users. Matching based on multiple dimensions (such as topics based on reviews and categories based on temporal check-ins) has been shown to dramatically increase the accuracy of POI matching [3].

In summary, this position paper takes the stance that *Spatial Search* is not only about *spatial*. In many cases dimensions outside of X, Y and Z can and should be employed to enhance spatial search. A discussion concerning the role that other dimensions play in searching inherently spatial data would be of considerable interest from my perspective.

References

[1] 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

[2] Raubal, M., Miller, H. J., & Bridwell, S. (2004). User-Centred Time Geography for Location-Based Services. *Geografiska Annaler: Series B, Human Geography 86*(4): 245–265.

[3] McKenzie, G., Janowicz, K., Adams, B. (2014) A Weighted Multi-Attribute Method for Matching User-Generated Points of Interest. *Cartography and Geographic Information Science* 41(2): 125–137, Taylor & Francis.

On Language Models for Places from User-Generated Content

VANESSA MURDOCK
Principal Applied Researcher
Microsoft
Email: vanmur@microsoft.com

We inhabit two worlds. We have offline interactions, which we have been developing and fine-tuning since the beginning of humanity, and we have online interactions, in which we employ our tremendous human capacity for social behavior. There are many points of overlap between the two spheres, such as friends we cultivate both online and offline, events that happen offline that we experience primarily online,[1] or online interactions that have offline consequences.[2] Location is one point of contact between the online and the offline experience.

Because GPS-enabled smart phones are widely used, we have an unprecedented amount of data about where people are, and what they were doing and thinking at the time. We can leverage this data directly for specific tasks, such as identifying the location a photo was taken, or more obliquely as a background model of a user's geographic context, to be used in ranking, recommendation or prediction.

One approach to harnessing geo-tagged social media is to turn it into an information retrieval problem. We propose to segment the globe into "cells" and populate each cell with the social media artifacts emanating from that cell. We treat each cell as a "document" and estimate a term distribution from it. This document representation of places allows us to rank places given a query, or compute a similarity function between a short text and a place.

In our work we investigated several segmentation schemes (based on fixed grid cells, on zip code boundaries, and on dynamically-sized grid cells). We have investigated several approaches to estimating the term distribution associated with a place, and several different ranking functions. We discovered that if the boundary of a place is not important, dynamically sized grid cells produce the most accurate language models (Murdock, 2014). Regardless of the boundaries, if the term distributions are estimated from user-generated content, they are better estimated with the user frequency than the term frequency (O'Hare and Murdock, 2013). Dirichlet smoothing produces better results than Jelinek-Mercer smoothing for fixed-sized cells, because fixed-sized cells have a greater variance in the number of terms used to describe the place (O'Hare and Murdock, 2013). This does not hold for dynamically sized cells, because the cell size is adjusted according to the vocabulary so there is less variance in the document length (Murdock 2014).

---

[1] An example would be experiencing a sporting, or political event through live Twitter feeds of people who are there.

[2] An example is #gamergate in which female games developers were chased from their homes, and canceled speaking events because of online harassment. [Rawlinson, K. (2 September 2014) Gamers take a stand against misogyny after death threats. *BBC News*. Retrieved from http://www.bbc.com/news/technology/ November 2014.]

Each data source is useful,[3] but they differ in the type of information they convey. For example, Flickr[4] is especially useful for location modeling because the user is very often describing the location of the photo. A caveat of Flickr is that the photos are often associated with a particular user context, and the text associated with the photographs may not be meaningful outside of that context (van House, 2007). Users frequently apply the same, or nearly the same, tag set to large numbers of uploaded photos, although this is easily remedied by removing exact duplicates, and estimating term distributions with the user frequency, which counts one instance of a term per user per location, to avoid a single user dominating the text representation of a place.

Foursquare[5] allows users to check in at a pre-existing place, or create a new place. They can also enter a short description or recommendation of a place (a "tip") that other users can see. Foursquare focuses on business venues, although it is not limited to them.

While Flickr and Foursquare data link text to a place very directly, Twitter[6] does so to a lesser degree. People often tweet about events or places around them, but they frequently tweet about topics unrelated to their current location. Mobile queries are also less often related to the location of the user. However, both Twitter and mobile queries allow researchers to infer the places people habitually visit, which is helpful for determining, among other things, which places are popular (Liu, 2013), and the user's likely itinerary (de Choudhury et al. 2010).

A criticism of social media is that it has a low signal to noise ratio. For every gem in a Twitter feed, there are thousands of status updates that are not informative to users outside of the sender's social context (Naaman et al. 2010). The sheer volume of user-generated content offsets this. For example, while in Twitter a very small percentage of tweets are associated with geographical coordinates, it is still a vast number of tweets. Flickr currently has hundreds of millions geo-tagged images. Foursquare users enable location by default, and so the majority of the Foursquare data is geo-tagged. Even queries from search engines on mobile devices are annotated with the coordinates of the users, in enormous volumes.

Social media and other user-generated content have been shown to be very useful for a variety of prediction, ranking and recommendation tasks. Among them, identifying the user's location (Kinsella 2010 et al., Cheng et al. 2010), disambiguating points of interest (Rae et al. 2012), improving the ranking of search engine results (Bennett et al. 2011), inferring tourist itineraries (de Choudhury et al. 2010), and event recommendation (Quercia et al. 2010). It is a vast source of data, and growing. The challenge is to find creative ways to use it while respecting the user's privacy.

References

P. N. Bennett, F. Radlinski, R. W. White, and E. Yilmaz (2011). Inferring and Using Location Metadata to Personalize Web Search. Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 134–144.

---

[3] Flickr, Twitter and Foursquare have public APIs granting limited access to their data.

[4] https://www.flickr.com visited November 2014.

[5] https://www.foursquare.com visited November 2014.

[6] https://www.twitter.com visited November 2014.

Z. Cheng, J. Caverlee, and K. Lee. You Are Where You Tweet: A Content-based Approach to Geo-Locating Twitter Users (2010). Proceedings of the 19th ACM International Conference on Information and Knowledge Management, pp. 759–768.

M. de Choudhury, M. Feldman, S. Amer-Yahia, N. Golbandi, R. Lempel, and C. Yu (2010). Automatic Construction of Travel Itineraries Using Social Breadcrumbs. Proceedings of the 21st ACM Conference on Hypertext and Hypermedia, pp. 35–44.

S. Kinsella, V. Murdock, and N. O'Hare. I'm Eating a Sandwich in Glasgow: Modeling Locations with Tweets (2011). Proceedings of the 3rd International Workshop on Search and Mining User-Generated Content, pp. 61–68.

B. Liu, Y. Fu, Z. Yao, and H. Xiong (2013). Learning Geographical Preferences for Point-Of-Interest Recommendation. Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 1043–1051.

V. Murdock. Dynamic Location Models (2014). Proceedings of the 37th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 1231–1234.

M. Naaman, J. Boase, and C-H. Lai (2010). Is it Really About Me? Message Content in Social Awareness Streams. Proceedings of the ACM Conference on Computer Supported Cooperative Work, pp. 189–192.

N. O'Hare and V. Murdock (2013). Modeling Locations with Social Media. *Journal of Information Retrieval* 16(1): 30–62.

A. Rae, V. Murdock, A. Popescu, and H. Bouchard (2012). Mining the Web for Points of Interest. Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 711–720.

D. Quercia, N. Lathia, F. Calabrese, G. di Lorenzo, and J. Crowcroft (2010). Recommending Social Events from Mobile Phone Location Data. Proceedings of the IEEE International Conference on Data Mining (ICDM), pp. 971–976.

N. van House (2007). Flickr and Public Image-Sharing: Distant Closeness and Photo Exhibition. CHI '07 Extended Abstracts on Human Factors in Computing Systems, pp. 2717–2722.

Visual Spatial Search

SHAWN NEWSAM
Associate Professor Electrical Engineering and Computer Science
Founding Faculty
University of California, Merced
Email: snewsam@ucmerced.edu

This white paper examines the interplay of visual information and location over the last two decades with a particular focus on spatial search. We describe how visual information can be a powerful spatial index for a range of search applications. We also track how advances in computer algorithms and technology have made visual spatial search an increasingly personal and egocentric experience to the point where our camera-equipped wearable devices are able to answer questions such as "Where am I?" and "What is that there?"

The growing collections of digital images in the 1990s motivated novel approaches for organizing and searching this data. Manual annotation is not scalable so methods to automatically index the visual content of the images were needed. Computer scientists working in multimedia proposed content-based image retrieval (CBIR) as a solution. CBIR skirts the challenge of semantic image understanding by leveraging low-level visual feature descriptors—color, texture, etc.—to search large sets of images for those most visually similar to a query image. Overhead images from satellite or airborne platforms are a natural application domain of CBIR and so I spent part of my Ph.D. studies at UC Santa Barbara developing content-based geographic image retrieval (CBGIR) systems as part of the NSF funded Alexandria Digital Library project. The figure on the right illustrates a query from an online CBGIR demo[1] where a user has selected a freeway region as the query and the system has returned images that are most similar with respect to color in the first row and most similar with respect to texture in the second row. CBGIR systems support visual spatial search by finding regions of the Earth that are similar, at least visually, to a region of interest. They can also be used to annotate novel images through label propagation.

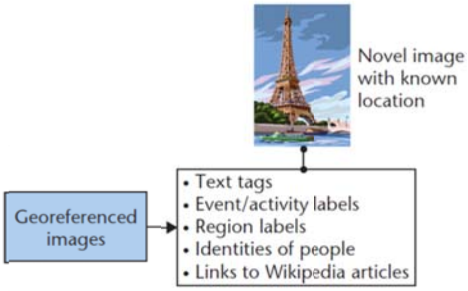

The advent of the web and GPS-equipped cameras, such as in smart phones, has made available a wealth of georeferenced social multimedia. In a recent position paper[2] [1], I postulate that this data is really a form of volunteered geographic information (VGI) since it supports geographic knowledge discovery through various forms of spatio-visual search. As depicted on the left, collections of

---

[1] http://vision.ucmerced.edu/GIR/

[2] References for the work described in the rest of this paragraph can be found in [1].

geo-referenced images, such as available at Flickr and similar photo sharing websites, have been used to semantically annotate novel images whose locations are known. Specifically, spatio- visual search has been used to suggest tags such as "surfer," "wave," and "Santa Barbara" for a photograph of someone surfing in Santa Barbara, California; for assigning a constrained set of event/activity labels such as "a visit to the beach" or "wedding"; for annotating groups of images at the event ("skiing") or scene("coast") level; for annotating the identities of people appearing in an image; and for linking images, such as a photograph of the Arc de Triomphe, to relevant Wikipedia articles. These collections have also been used to annotate geographic locations such as visually annotating prominent landmarks with representative images at the city and world-wide scales; to suggest representative tags and images for Geographic locations; and to automatically generate tourist maps showing popular landmarks as vectorized icons. These collections have also been used to learn new geographic information such as spatially varying (visual) cultural differences among concepts such as "wedding cake" and discover interesting properties about popular cities and landmarks such as the most photographed locations. In my own work, I have used these collections to map land cover [2], land use [3], and "scenicness" [4].

Large collections of georeferenced images have also been used to determine the locations of novel images as shown on the right. The general idea is to perform a visual search through a collection and propagate locations to the novel image usually in a probabilistic framework. The problem of geoferencing an image is also the focus of the Finder project of Intelligence Advanced Research Projects Activity (IARPA).

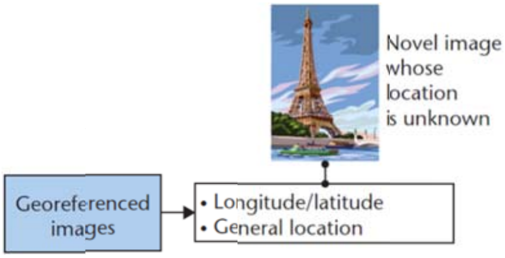

The Photo Tourism research project [5] uses state-of-the-art image matching to co-register largecollections of images of famous landmarks. These images are found simply through a keyword search ("Notre Dame Cathedral" at Flickr for example), and, while their geographic locations are not known, their relative locations can be estimated through the automatic extraction and matching of visual keypoints. On the left is shown the camera locations of several hundred Flickr images of Notre Dame Cathedral as determined through keypoint matching. Once registered, the images support intuitive interactive browsing. The matched keypoints can also be used to create a 3D point cloud of the object as depicted in the figure.

The age of wearable computers we are entering will enable novel applications of personal visual spatial search. Location constrained visual matching is already being used to overlay informative text in augmented reality systems to answer the question "What is that there?" Similar matching can locate users to answer the question "Where am I?" especially in GPSchallenged environments such as indoors. It is not far-off that these devices will perform personal simultaneous location and mapping, a well-studied problem in robotics where an agent uses sensors to learn and locate itself within its environment. This would allow us to leave virtual breadcrumbs for navigating out of unfamiliar environments. These devices will also support the most basic and personal of spatial searches such as "Where are my keys?" through a combination of recording and searching our past actions as well as analyzing the spatial layout of a scene in real-time. Real-time scene analysis could also help us efficiently locate things in complex environments such as a friend in a crowd or an item on a store shelf.

References

[1] S. Newsam, "Crowdsourcing what is where: Community-contributed photos as volunteered geographic information," *IEEE Multimedia: Issue on Mining Community-Contributed Multimedia*, 17(4), pp. 36–45, 2010.

[2] D. Leung and S. Newsam, "Proximate sensing: Inferring what-is-where from georeferenced photo collections," *IEEE International Conference on Computer Vision and Pattern Recognition*, 2010.

[3] D. Leung and S. Newsam, "Exploring geotagged images for land-use classification, "*ACM International Conference on Multimedia: Workshop on Geotagging and Its Applications in Multimedia*, 2012.

[4] L. Xie and S. Newsam, "IM2MAP: Deriving maps from georeferenced community contributed photo collections," *ACM International Conference on Multimedia: Workshop on Social Media*, 2011.

[5] N. Snavely, S. M. Seitz, R. Szeliski, "Photo tourism: Exploring photo collections in 3D," *ACM Transactions on Graphics (SIGGRAPH Proceedings)*, 25(3), pp. 835–846, 2006.

Image-Based Spatial Search for Telecollaboration

BENJAMIN NUERNBERGER
Department of Computer Science
University of California, Santa Barbara
Email: bnuernberger@cs.ucsb.edu

With

MATTHEW TURK
Department of Computer Science
University of California, Santa Barbara
Email: marcus.ufv@gmail.com

TOBIAS HÖLLERER
Department of Computer Science
University of California, Santa Barbara
Email: marcus.ufv@gmail.com

Searching for places or objects in a collection of photographs that are spatially related is a task that is becoming more and more relevant in today's digital world. Photographers, designers, and hobbyists must often search through thousands of photos to find what they are looking for. People search for places visually through applications such as Google Street View or Microsoft PhotoSynth. Furthermore, collaborators are increasingly using images for integrating the spatial world into remote collaboration. In this position paper, we approach spatial search from the viewpoint of virtually navigating a remote world represented by a set of spatially related images for the ultimate purpose of improved telecollaboration. Telecollaboration is an enhanced form of videoconferencing that allows participants to collaborate in more ways than just seeing and hearing each other. We first describe related work in image-based spatial search, then present our recent research in applying image-based spatial search to telecollaboration, and finally give our thoughts on some current challenges and directions for visual spatial search. Searching for images associated with physical locations has been approached in various ways in the past. Using solely automatic GPS annotations and manually added keywords has long been the naïve approach to searching through such images that are tied to a spatial location. More recently, 2D image features (such as SIFT features) and visual vocabularies have been used to create systems that match locations or objects across videos or photographs [1]. In addition, research in structure from motion and stereo vision (and also vision with depth sensors) has allowed computers to analyze spatially related photos to recover the 3D relation between them and the 3D structure of the commonly observed scene [2]. Researchers have also been analyzing image-based localization—that is, determining the geo-location of a single photograph; this is useful for mobile localization when GPS is often unavailable [3].

Recently we in the Four Eyes Lab at UCSB have been researching ways to integrate the physical environment into telecollaboration [4]. In our case, a local user holds a tablet device which transmits a video stream to a remote user. The remote user is then able to spatially annotate the local user's scene by using augmented reality. Furthermore, the remote user can virtually navigate or search through the scene by using a subset of the video stream frames as "key-frames" that encompass different viewpoints of the scene's spatial area. By using augmented reality and image-based navigation, our prototype telecollaboration application allows users to virtually annotate the remote physical world and to search through it in a visual way. Our results have shown that users prefer our user interface over more traditional telecollaboration interfaces.

In this context, we have been exploring the challenges of virtual navigation of a remote scene for telecollaboration. Virtual navigation typically encompasses three tasks—exploration, search, and inspection. To investigate how different user interfaces affect user performance, we used a naïve search task in an exploratory user study and found insights into designing future user interfaces for image-based navigation of virtual environments [5]. Later, we expanded our interface to use a touchscreen and found that users prefer our interface over others [6].

By investigating virtual navigation for telecollaboration, we have unearthed aspects and challenges of visual spatial search that have not been explored much in the past. This includes the real-time collaborative aspect of visual spatial search. Previous systems, such as Google Street View or Microsoft PhotoSynth, must build offline a virtual scene model for users to spatially search through; our research, however, brings in a real-time element to building this model through which users can spatially search for items of interest.

Another challenge for visual spatial search in the context of telecollaboration is the cognitive aspects of designing the user interface. In constrained situations, such as Google Street View, clicking in predefined areas (e.g., roads) makes it easy to search and navigate through images with relatively low cognitive load on the user. In unconstrained situations, however, current visualizations for image-based navigation interfaces are not sufficient and cause high cognitive load on the user [5]. On the one hand, users can simply scroll through lists of photos in a brute force manner with low cognitive load but with also extremely low throughput. On the other hand, the system may be able to index visual features common to multiple images [1] or utilize 3D structure [4] to increase search throughput while maintaining a low cognitive load.

Perhaps the two main challenges currently for visual spatial search are what some researchers call "perceptual aliasing"—when very similar visual information appears in multiple places, making it hard to distinguish between them—and "image variability"—when the visual appearance of a certain place changes temporally, making it hard to recognize the same place over time. For example, it is hard to geo-localize a photo of a bicycle without any contextual information and it may be hard to recognize the same park from photos taken during different seasons, such as winter and summer. However, for small areas, such as might occur in telecollaboration, identifying the relative associations between places may be more achievable. All of these considerations lead us to interesting research questions, such as, "What is the relationship between searching through a collection of images and navigating/exploring a remote spatial environment?" And, "Is it possible to juxtapose research advances in image-based search and 3D scene reconstruction, along with human computer interaction techniques, to enable efficient and intuitive visual spatial search?"

In conclusion, we have begun to explore visual spatial search in the context of telecollaboration, and this has opened up important questions in how to apply previous research advances in visual search to the current challenges of realizing visual spatial search. As research continues to progress towards addressing the spatial component of visual search, we will begin to see systems that enable humans to more intuitively and efficiently search through both the visual and spatial realm.

Works Cited

[1] J. Sivic and A. Zisserman, Video Google: A text retrieval approach to object matching in videos, in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2003.

[2] N. Snavely, S. M. Seitz and R. Szeliski, Photo tourism: exploring photo collections in 3D, in *ACM Transactions on Graphics (TOG)*, 2006.

[3] T. Sattler, B. Leibe and L. Kobbelt, Fast Image-Based Localization using Direct 2D-to-3D Matching, in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2011.

[4] S. Gauglitz, B. Nuernberger, M. Turk and T. Hollerer, World-Stabilized Annotations and Virtual Scene Navigation for Remote Collaboration, in *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST)*, 2014.

[5] B. Nuernberger, S. Gauglitz, T. Hollerer and M. Turk, Poster: Investigating Viewpoint Visualizations for Click & Go Navigation, in *Proceedings of the IEEE Symposium on 3D User Interfaces (3DUI)*, 2014.

[6] S. Gauglitz, B. Nuernberger, M. Turk and T. Hollerer, In Touch with the Remote World: Remote Collaboration with Augmented Reality Drawings and Virtual Navigation, in *Proceedings of the ACM Symposium on Virtual Reality Software and Technology (VRST)*, 2014.

Spatial Search—
The Next Evolutionary Step for the Post-SDI Age

ED PARSONS
Geospatial Technologist
Google Research
Email: eparsons@google.com

In this position paper I aim to bring an industry perspective to the problem of the discovery and access to information with a geospatial context, both the traditional representation of entities within information systems and the emerging field of digital interactions with the entities themselves, the physical web or internet of things. The industry perspective of course is not one of unanimity, indeed one could argue there are clear dividing lines between the "professional" data producer dominated world of Spatial Data Infrastructure (SDI) developers and the world of mass market web companies and app developers for whom the perceived complexity of metadata standards and service oriented architectures has been ignored in favor of rapidly iterated systems.

The Dominance of the Search Metaphor

The fact that "to google" has become a verb is the direct result of the popularity and dominance of search engines including, Bing and Yahoo as interfaces to find information online. However the widespread adoption of search goes far beyond web search engines, as users we now expect our email clients, set top boxes and even station ticket machines to offer the ability to search for information rather than follow some arbitrary menu system or sequential series of dialogues.

In the past few years the search metaphor has been further developed as a result of technological developments in voice recognition, natural language processing and artificial intelligence resulting in conversational search where a user talks to their mobile device which responds with results and allows further contextual refinement of these results in the form a natural conversation.

The effort to make the user interface to an information system in essence disappear has been a long sought after and is almost a reality. Almost as a mirror to the simplicity of the user interface the amount of processing required and the volumes of structured information required to allow conversational search is immense, the knowledge graph the internal repository of structured data used by Google contains over a billion discrete facts.

The persistence of the Map Chest Metaphor

In contrast to the rapid development is web search, the geospatial community has continued to follow a slower development path embracing broader policy and legal issues as well as technological under the umbrella term of Spatial Data Infrastructure (SDI). The SDI approach is characterized by the clear separation of information and its metadata, following the traditional map library approach of a card catalogue and map chests.

In contrast to the web search approach where the creation of semantic understanding comes from machine learning, the creation of metadata remains a largely manual process despite years of tool development. Another contrast is that as a result of this separation it often requires separate processes and tools to firstly interrogate metadata catalogues and then as a secondary step potentially access the information itself.

As a result of this architecture SDI's and the information they contain and remained solely within the GIS and GI Science community, indeed in most cases the information published a local, regional and national scales via SDI portals remains opaque to web search as it is published not using entity level uniform resource locators (URL's) but via obscure domain specific web services.

Time to mix metaphors—a call to action!

I argue that for the maximum benefit to be achieved from publishing geospatial datasets the GIS and GI Science community must adopt the paradigm of web search and publish geospatial content using mainstream web standards, specifically to follow a modified version of principles established by Tim Berners-Lee on publishing linked data.

Berners-Lee proposes that data should be published using open (without restriction) licenses, using machine readable and non-proprietary standards. I would suggest perhaps some more lightweight methods of encoding semantics rather than the suggested RDF approach using simpler vocabularies such as those suggested by schema.org based on microdata.

There is still a considerable research problem of course in the development of geospatial semantic knowledge and its creation at scale, fuzzy geographical concepts such as "near," the hierarchy of places from a societal perspective and issues of scale and time all require new techniques of machine learning to harvest from existing sources of geospatial content online.

In the near future the web of information representing real world objects will be joined by the web of the objects themselves, the internet of things or the physical web where the objects around us become locations on the web as well as "real world" locations offer a new perspective on the problem of geospatial search.

Ultimately will we be able to answer the follow indicative questions . . . "Will I need a jacket when I arrive Monday?" and "Is the car2go parked close to a shop where I can buy a new phone charger?"

Enabling Flexible Search for ArcGIS Online Using Semantic Web Technologies

SATHYA PRASAD
Applied Research Architect
Applications Prototype Lab
Esri
Email: sprasad@esri.com

ArcGIS Online (http://www.arcgis.com/) is a Web platform developed by Environment System Research Institute (ESRI). As a worldwide geoportal, it enables GIS users from different countries to create, edit, and upload geo-data, Web maps, services, web applications and GIS tools, as well as to share these GIS resources with colleagues, friends, and the general public. ArcGIS Online resources (including all the maps, data, and resources) are called *items*, and they are accompanied by a rich set of metadata, including information about title, snippet, description, user, user groups, thumbnail image, creation date, modification date, comments, ratings, and so forth. Based on these metadata, users can browse through the collection of items, or sort them by features such as the popularity or themes.

In order to efficiently find the right items in the ocean of ArcGIS Online data, a RESTful Web API has been developed which enables search and retrieval of the items along with the users and groups. However, these existing APIs also restrict the way that the ArcGIS Online resources can be searched. For example, if the users would like to understand the information about both an item and the user who have created this item, then two queries have to be sent out separately to get the corresponding information. In some other cases, a user may want to know the information about the layers contained within a map, but such information cannot be easily searched simply because no corresponding API functions are there. These issues emerge because the developed APIs restricts the search capabilities, and therefore we cannot search the items when no corresponding APIs have been built even if there exist items for the users' queries.

While it is always possible to expand the API functionalities by adding more code, users in different organizations and domains may easily come up with specific queries, which are beyond the expectation of the developers. Consequently, our goal in this position paper is to enable flexible search based on the existing ArcGIS Online search architecture. By "flexible search," we mean that our system should allow users to input any queries without being restricted by the capabilities of the pre-designed APIs. While this work has been experimental with ArcGIS Online, the proposed methods can also be applied to other geoportals (e.g., Data.gov).

The solution we propose in this position paper is based on the Semantic Web technologies and the principles of Linked Data. Semantic Web is the third generation of the Web proposed by Tim Berners-Lee, the father of the World Wide Web. Semantic Web employs a set of standards for organization, publishing, and retrieving data from the Web. Resource Description Framework (RDF) is a standard data model used on Semantic Web to store data, and it organizes data into subjects, predicates, and objects. The term "Linked Data" can be used to represent two meanings. On one side, it refers to the four standard principles to organize data, and on the other side, it can refer to the data that gets organized by following the four principles.

The idea we have in this paper is to organize ArcGIS Online metadata into Linked Data that can be represented as a graph of nodes and edges. Then, users can query items by entering this graph from any node and in any direction. Thus, instead of adding a huge number of API functionalities, we organize the data in a "smart" way, that can answer any query from the users as long as the corresponding items exist.

To test this idea, we experimented with a sample of ArcGIS Online data. Ontologies have been defined for items on ArcGIS Online, which can be primarily divided into classes of items, item types, users, and groups. The ontologies are not designed from scratch but based on the schema ArcGIS Online already has in the JSON response from the REST API. Therefore, developers who are familiar with the REST API can learn how to query the RDF data more quickly. Using Java and Jena API, we have designed an application called "RDFConverter" to convert data from its original style to RDF. We then published these data following the Linked Data principles onto a triple store called Parliament. We choose Parliament instead of other triple stores as the SPARQL endpoint because it has implemented GeoSPARQL, which we are very interested to investigate. After publishing the data, we built a front-end web application which directly talks to the SPARQL endpoint and consumes RDF data. The implemented prototype portal supports the flexible queries without having to design corresponding API functions. For example, a quick SPARQL search can find out the owners of the maps that use a particular basemap. We can also find out the top 10 most popular basemaps based on their usage. In addition, we can search into the map and find information about layers (e.g., a Web map with layers about "population" and "disaster"). These are some of the examples, and all these queries can be complete within a single request.

To sum up, this work employs Semantic Web technologies for ArcGIS Online to enable a more flexible search. Developing "smart data" instead of only "smart applications" (i.e., adding many API functions) could be an important direction for enhancing spatial search.

Geographic Information Retrieval: Are We Making Progress?

Ross Purves
Department of Geography
University of Zurich
Email: rspurves@gmail.com

What is spatial search? And perhaps more importantly, what are the implications of developing systems which allow some form of spatial search? In this position paper I aim to explore these questions from the perspective of Geographic Information Retrieval, a term originally coined by Larson (1996) in his seminal paper in 1996 and, refined to more specifically refer to unstructured information, primarily as found in text documents, in an editorial written by Chris Jones and myself in a special issue of IJGIS published in 2008 (Jones and Purves, 2008). In that editorial we set out what we believed to be key challenges for GIR. These challenges were based on our experiences in developing the SPIRIT system, one of a number of early GIR systems (e.g., Chen et al., 2006; Lieberman et al., 2007) which sought to bring some form of geographic intelligence to information retrieval. As such, the challenges closely mimicked what believed to be the necessary components of such a system, and put little emphasis on, rightly or wrongly, developments in areas such as linked data and the semantic web.

In the following, I will briefly list these challenges, and attempt to explore progress by looking at the literature from two viewpoints. Firstly, I am interested in the broad range of methodological and disciplinary approaches applied to the different challenges. Does lack of attention to a challenge imply that we simply identified the wrong questions, or that progress in other areas is required before such challenges can be met? Secondly, since spatial search and GIR have now become mainstream in computer science, I am interested in what useful ideas, if any, GIScience and geography might have to offer in this research field.

Challenges and Progress

Detecting geographical references in the form of place names and associated spatial natural language qualifiers within text documents and in users' queries

When we wrote our editorial, the focus of most work in GIR was on text documents published on the web, of variable lengths, and often containing rich geographic information. Extracting locations from documents had been rightly identified as an important challenge, and the focus here was not *per se* on what is known as toponym resolution (identifying a single, unique referent for a given place name), but rather toponym recognition (identifying whether a token or phrase within a document can be treated as a spatial referent) (e.g., Leidner, 2008). This area of research had already been the subject of attention from computer linguistics as part of the more general task of Named Entity Recognition. Queries can be seen as forerunners to the processing of other very short documents, such as Tweets, as the advent of georeferenced feeds has provided a perfect opportunity to experiment with, for example, approaches based on machine learning (e.g., Roller et al., 2012). Indeed, many approaches based on machine learning conflate toponym recognition and resolution and remove the need for both gazetteers and heuristics based on the ways in which place names are used. GIScience has seen increasing attempts to use corpora such as Twitter as sources of information about spatial language— however the extent to which, for instance, some of the ideas of naïve geography (Egenhofer and Mark, 1995) and the related use of spatial language has informed this research appears limited.

Disambiguating place names to determine which particular instance of a name is intended

Having identified candidate referents in documents, a key task is their disambiguation, and thus toponym resolution. Humans typically excel at this, by using both clues within the document and other contextual information, to the extent that booking a flight in error to Sydney, Nova Scotia rather than Australia is the subject of a news story (http://news.bbc.co.uk/2/hi/uk/2172858.stm). However, this task remains a challenging one for automated systems, which typically adopt relatively simple approaches, such as using ancillary information mined from other content (e.g., population counts or co-occurrence with other terms) or simple geometric measures to calculate distances to the nearest unambiguous referent identified (Smith and Crane, 2001; Buscaldi, 2011). However, many of these approaches assume that geography is random (i.e., that ambiguous toponyms have no spatial autocorrelation) despite clear evidence that this is not always the case (Brunner and Purves, 2008). Furthermore, most approaches focus on toponyms related to settlements (Leidner, 2008) despite the importance of dealing with more fine-grained toponyms referring to, for example, individual locations along the route of a hike.

Geometric interpretation of the meaning of vague place names, (e.g., "Midlands") and of vague spatial language (e.g., "near")

Many toponyms are not found in official gazetteers. Often, these toponyms are also vague, and in fact seemingly well-defined toponyms may be being used in natural language in a vague sense. In addition, the use of spatial language can take seemingly well-defined concepts and render them vague. Such concepts have moved from being interesting ideas for GIScientists to consider (e.g., Montello et al. 2003; Jones et al. 2008a) to being of fundamental importance in resolving queries successfully. If around 13% of queries really are in some way spatial (Jones et al., 2008b), it is likely that a sizable portion of these will refer to either vague (and vernacular) place names as well as vague spatial language. Thus, considerable effort has been expended on identifying the geometric regions associated with such place names (e.g., Grothe and Schaab, 2009; Keßler et al., 2009) though often the notion of vagueness has been discarded along the way (e.g., Jones et al., 2008a). Indeed, typical search engine representations of near seem to simply assume that metric distance is a "good enough" method for ranking, and base this on point-based models of geographic entities. Finding place names which are not already in gazetteers, and associating these with either metric or topological models of space is an area of research with much potential, especially if it can be scaled up to be relevant for large digital corpora.

Indexing documents with respect to geographic context and non-spatial thematic content

The notion of indexing is central to all efficient search. Early research in GIR explored the influence of different approaches to index construction (e.g., text vs. space-primary) (Vaid et al., 2005) though in practice often simple approaches based on multiple indexes and intersection operations appear to have been effective. However, three other important questions can be posed with respect to indexing for spatial search. (1) What should be indexed— should a document be represented as a "bag of points" corresponding to all resolved locations found in the text, as a single footprint, or through some other intermediate representation. (2) The approach taken to spatial indexing (e.g., space or object-directed indexing) has important implications for the questions that can then be explored through the corpus, beyond basic retrieval operations. Such questions are also where potential benefits for geography lie; for example if we wish to map corpora onto existing spatial data, then space-directed indexes ensure that a complete tessellation of space, such that all locations can be linked to content at some given (not necessarily equivalent) granularity. (3) It may be that queries and documents could usefully be represented in different ways. Even where document footprints are stored as simple points, it makes sense to use the observation "topology defines, metric refines" and allow query footprints to take more complex forms, enabling for example containment queries to be meaningfully formulated.

Ranking the relevance of documents with respect to geography as well as theme

Relevance of documents with respect to spatial search has two obvious components – thematic and geographic relevance. Producing a meaningfully ranked list of documents therefore requires that both of these components are considered. Many commercial search engines appear to treat these as independent variables, displaying documents which are thematically relevant on a map, with geographic relevance either being treated as binary variable, or sometimes ranked according to distance. Kreveld et al. (2005) explored how geographic and thematic relevance could be considered with respect to some ideal score, effectively representing the two relevance dimensions as orthogonal. A further important aspect of ranking is diversity— in the case of spatial search it makes sense to consider how cluster documents are, and attempt to present the user with not only relevant, but geographically (and thematically?) relevant results (Tang et al., 2010). It is also possible to consider incorporating more context into producing ranked sets of results, in particular with respect to mobile search, where user behaviour may allow, for example, predictions to made about future movement (Mountain and Mcfarlane, 2007) or relevance ranking methods which more fully take into account geographic context in a holistic way (Cai et al., 2007; De Sabbata and Reichenbacher, 2012). Interestingly, there seems to be little formal cross-over between work on spatial cognition and relevance ranking in GIR, despite the obvious potential of considering, for example, notions of salience in developing relevance ranking methods.

Developing effective user interfaces that help users to find what they want

User interfaces for geographic search have developed little beyond the original prototypes of SPIRIT in the early 2000s. Query formulation generally supposes a triplet of the form <theme><spatial relationship><location> which can be achieved either by use of a simple structured interface, a multi-modal interface (using a query box to specify theme, the map extent to imply location and typically ignoring spatial relationships), or by processing free text entries to extract some or all of the components of the above triplet. Equally, and perhaps even more surprisingly, results display in spatial search has not progresses much beyond the display of points on maps, despite some attempts to develop more novel query display approaches (e.g., Hobona et al., 2006; Carmo et al. 2007). Again, here there is little evidence of meaningful crossover from the geovizualization community directly dealing with search, as opposed to exploration, tasks.

Developing methods to evaluate the success of GIR

When we wrote our editorial the importance of developing effective approaches to evaluation was very clear for many researchers, and indeed the GeoCLEF campaign (e.g., Mandl et al., 2008) focussed on a system-centred strategy to allow comparison between systems. However, this campaign was largely unsuccessful in demonstrating the benefits of GIR over and above simple tweaks to textual systems. However, evaluations of individual GIR systems have often shown benefits over simple textual baselines. The obvious question that one can then pose is why? I believe a number of aspects particular to GIR have been neglected here. Firstly, corpora typically have very specific properties, with for example target audiences having a strong influence on the spatial language, and in particular the toponyms, used. Secondly, evaluating geographic relevance, especially at a local scale appears to be much more challenging than thematic (e.g., Clough et al., 2006, Ostermann et al., 2013), often requiring detailed local knowledge. Traditional and emerging approaches to evaluation based on pooled judgements and crowd sourcing are ill-suited to judging such nuances, and the need for more qualitative user-centred evaluation for judging long-tail queries should not be underestimated.

A second key point in developing effective evaluation strategies concerns harnessing the information which can be obtained from query logs in order to better define how users actually interact with systems. Work by Jones et al. (2008b), importantly performed under the auspices of a search engine company, suggested considerable potential for such explorations. The controversial release of the AOL query logs also demonstrated the potential for analysing the properties of geographic queries in a traditional search engine (Henrich and Lüdecke, 2007), however outside the commercial world access to such data remains rare. Better understanding how users formulate geographic queries, and linking this to basic notions of spatial cognition seems to me to be an obvious area where interdisciplinary research could rapidly bring tangible benefits.

Developments

Perhaps the most obvious development which could influence spatial search is the parallel rise of mobile devices and the use of social media. Mobile devices have important implications for context and the way it is used in geographic search, but most applications have continued to focus on the use of well structured (e.g., point of interest) data, and the integration of unstructured data beyond mobile versions of traditional search mechanisms has seen surprisingly little development. The rise of social media has resulted in a need for methods to index and query what are effectively very short documents, and approaches based on machine learning dominate here. A further, perhaps less obvious development is the increase in the open availability of large corpora of legacy documents (e.g., http://chroniclingamerica.loc.gov/ or http://textberg.ch/site/en/welcome/) often with a specific geographic or thematic focus. The potential of such corpora for answering geographic questions has been, in my view, neglected, and I believe some of the greatest potential for interdisciplinary collaborations between researchers.

In looking back at the challenges we set, a few points emerge. Toponym recognition and resolution remain absolutely fundamental, and despite Leidner's pleas (Leidner, 2006), gold standards allowing us to evaluate and compare methods have not emerged. Indexes, though interesting, have seen much less research than some areas, at least in part because simple approaches appear to often be "good enough." By contrast, dealing with vague and vernacular placenames was, and remains, a hot topic, despite the fact that clear evidence of concrete benefits for retrieval task (as opposed to assertions) perhaps remain illusory. Relevance ranking has remained an area which has received relatively little attention, despite the seemingly clear importance of presenting results which are spatially divergent, and indeed the presentation of such results has obvious links with the interfaces of systems providing GIR. Finally, evaluation, despite being the subject of significant attention through the GeoCLEF series remains an area with much potential, where we can perhaps more effectively learn and pool research results.

Acknowledgements: This is a position paper, and the views asserted here are mine and mine alone. Nonetheless, these views are the result of a long and valuable collaboration with many colleagues, and I would like to acknowledge in particular my discussions with Paul Clough, Curdin Derungs, Alistair Edwardes, Chris Jones and Mark Sanderson.

References

Brunner, Tobias Josef, and Ross Stuart Purves. (2008) Spatial autocorrelation and toponym ambiguity. In *Proceedings of the 2nd international workshop on Geographic information retrieval*, pp. 25–26. ACM.

Buscaldi, Davide. (2011) Approaches to disambiguating toponyms. *SIGSPATIAL Special* 3(2): 16–19.

Carmo, Maria Beatriz, Ana Paula Afonso, and Paulo Pombinho Matos. (2007) Visualization of geographic query results for small screen devices. In *Proceedings of the 4th ACM workshop on Geographical information retrieval*, pp. 63–64. ACM.

Chen, Yen-Yu, Torsten Suel, and Alexander Markowetz. (2006) Efficient query processing in geographic web search engines. In *Proceedings of the 2006 ACM SIGMOD international conference on Management of data*, pp. 277–288. ACM.

Clough, Paul D., Hideo Joho, and Ross Purves. (2006) Judging the spatial relevance of documents for GIR. In *Advances in Information Retrieval*, pp. 548–552. Springer Berlin Heidelberg.

Clough, Paul D., Hideo Joho, and Ross Purves. (2006) Judging the spatial relevance of documents for GIR. In *Advances in Information Retrieval*, pp. 548–552. Springer Berlin Heidelberg.

De Sabbata, Stefano, and Tumasch Reichenbacher. (2012) Criteria of geographic relevance: an experimental study. *International Journal of Geographical Information Science* 26(8): 1495–1520.

Egenhofer, Max J., and David M. Mark. (1995) *Naive geography*. Springer Berlin Heidelberg.

Grothe, Christian, and Jochen Schaab. Automated footprint generation from geotags with kernel density estimation and support vector machines. (2009) *Spatial Cognition and Computation* 9(3): 195–211.

Henrich, Andreas, and Volker Luedecke. (2007) Characteristics of geographic information needs. In *Proceedings of the 4th ACM workshop on Geographical information retrieval*, pp. 1–6. ACM.

Hobona, Gobe, Philip James, and David Fairbairn. (2006) Multidimensional visualisation of degrees of relevance of geographic data. *International Journal of Geographical Information Science* 20(05): 469–490.

Jones, Christopher B., Ross S. Purves, Paul D. Clough, and Hideo Joho. (2008a) Modelling vague places with knowledge from the Web. *International Journal of Geographical Information Science* 22(10): 1045–1065.

Jones, Rosie, Wei V. Zhang, Benjamin Rey, Pradhuman Jhala, and Eugene Stipp. (2008b) Geographic intention and modification in web search." *International Journal of Geographical Information Science* 22(3): 229–246.

Keßler, Carsten, Patrick Maué, Jan Torben Heuer, and Thomas Bartoschek. (2009) Bottom-up gazetteers: Learning from the implicit semantics of geotags. In*GeoSpatial semantics*, pp. 83–102. Springer Berlin Heidelberg.

Larson, R. R. (1996). Geographic information retrieval and spatial browsing.

Leidner, Jochen L. (2006) An evaluation dataset for the toponym resolution task. *Computers, Environment and Urban Systems* 30(4): 400–417.

Leidner, Jochen L. (2008) Toponym resolution in text: Annotation, evaluation and applications of spatial grounding of place names. Universal-Publishers.

Lieberman, Michael D., Hanan Samet, Jagan Sankaranarayanan, and Jon Sperling. (2007) STEWARD: architecture of a spatio-textual search engine. In*Proceedings of the 15th annual ACM international symposium on Advances in geographic information systems*, p. 25. ACM.

Mandl, Thomas, Paula Carvalho, Giorgio Maria Di Nunzio, Fredric Gey, Ray R. Larson, Diana Santos, and Christa Womser-Hacker. GeoCLEF 2008: the CLEF 2008 cross-language geographic information retrieval track overview. (2009) In *Evaluating Systems for Multilingual and Multimodal Information Access*, pp. 808–821. Springer Berlin Heidelberg.

Montello, Daniel R., Michael F. Goodchild, Jonathon Gottsegen, and Peter Fohl. (2003) Where's downtown?: Behavioral methods for determining referents of vague spatial queries. *Spatial Cognition and Computation* 3(2–3): 185–204.

Mountain, David, and Andrew Mcfarlane. (2007) Geographic information retrieval in a mobile environment: evaluating the needs of mobile individuals. *Journal of Information Science*.

Ostermann, Frank O., Martin Tomko, and Ross Purves. (2013) User evaluation of automatically generated keywords and toponyms for geo-referenced images. *Journal of the American Society for Information Science and Technology* 64(3): 480–499.

Purves, Ross S., Paul Clough, Christopher B. Jones, Avi Arampatzis, Benedicte Bucher, David Finch, Gaihua Fu et al. (2007) The design and implementation of SPIRIT: A spatially aware search engine for information retrieval on the Internet. *International Journal of Geographical Information Science* 21(7): 717–745.

Roller, Stephen, Michael Speriosu, Sarat Rallapalli, Benjamin Wing, and Jason Baldridge. (2012) Supervised text-based geolocation using language models on an adaptive grid. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pp. 1500–1510. Association for Computational Linguistics.

Smith, David A., and Gregory Crane. (2001) Disambiguating geographic names in a historical digital library. In *Research and Advanced Technology for Digital Libraries*, pp. 127–136. Springer Berlin Heidelberg.

Tang, Jiayu, and Mark Sanderson. (2010) Evaluation and user preference study on spatial diversity. In *Advances in Information Retrieval*, pp. 179–190. Springer Berlin Heidelberg.

Vaid, Subodh, Christopher B. Jones, Hideo Joho, and Mark Sanderson. (2005) Spatio-textual indexing for geographical search on the web. In *Advances in Spatial and Temporal Databases*, pp. 218–235. Springer Berlin Heidelberg.

Van Kreveld, Marc, Iris Reinbacher, Avi Arampatzis, and Roelof Van Zwol. (2005) Multi-Dimensional Scattered Ranking Methods for Geographic Information Retrieval.* *GeoInformatica* 9(1): 61–84.

Yu, Bo, and Guoray Cai. (2007) A query-aware document ranking method for geographic information retrieval. In *Proceedings of the 4th ACM workshop on Geographical information retrieval*, pp. 49–54. ACM.

Limiting Search with Simple Heuristics

LAEL SCHOOLER
Professor, Department of Psychology
Syracuse University
Email: lael.schooler@gmail.com

Simple heuristics are decision strategies that use limited information to make effective decisions in an uncertain world (Gigerenzer and Brighton, 2009; Katsikopoulos, Schooler & Hertwig, 2010). As a quintessential simple heuristic, Fast and Frugal Trees (FFTs) classify objects into one of two categories by asking a series of questions about the objects' characteristics (e.g., Luan, Schooler, and Gigerenzer, 2011; Martignon, Katsikopoulos, and Woike, 2008). Figure 1 shows a FFT that decides whether to prescribe macrolide antibiotics to children at risk for pneumonia (Fischer et al., 2002). The tree starts by asking whether the child has had a fever for more than 2 days. If the answer is "no," a decision is made (do not prescribe antibiotics) and no other characteristics are considered. If the answer is "yes," the tree asks the next question: Is the child older than three? If the answer is "yes," the decision is to prescribe antibiotics. If the answer is "no," then no antibiotics are given. Fischer and colleagues suggest that using their tree could substantially reduce unnecessary prescriptions of antibiotics. Beyond being both effective and efficient, FFTs are robust in that they can be trained on comparatively small samples; yet they still provide reliable out-of-sample predictions. But perhaps more importantly, they are easy to communicate. FFTs have successfully been developed in a variety of applied domains, including medical, legal, and financial decision-making (see Luan et al., 2011, for examples).

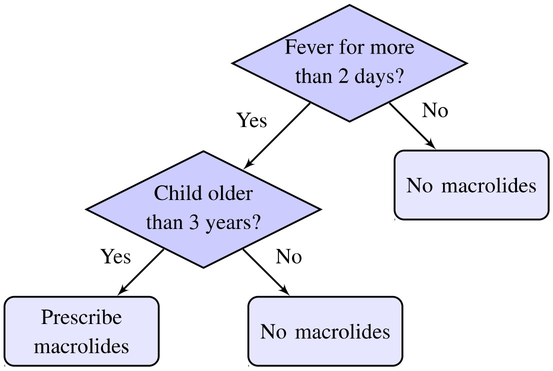

Figure 1: A FFT that decides whether or not to prescribe macrolide antibiotics to children at risk for pneumonia (Fischer et al., 2002).

A second class of simple heuristics make inferences about which of two options will score higher on a criterion of interest (e.g., will it be faster to take the left or right fork in the road?). A recent a sequential model for paired comparison, called *Δ-Inference*, stops information search when the difference between corresponding characteristics of the two options exceeds an aspiration level Δ (Luan, Schooler, and Gigerenzer, 2104). The accuracy of *Δ-Inference* peaks at values of Δ that on average are close to zero in both simulated and real task environments. In practice a Δ close to 0 means that *Δ-Inference* rarely considers more than one characteristic. Despite depending on limited search, *Δ-Inference* can achieve levels of accuracy as high as benchmark statistical models, such as linear regression. Both FFT's and *Δ-Inference* illustrate that when making inferences in the real world, it can pay to ignore information. Their success, however, depends heavily on the structure of the task environment. That is, the correlational structure of the objects and their characteristics. The good news is that many natural task environments have just the right structure that these heuristics need.

Some work within the simple heuristics framework has investigated spatial search. For instance, Hutchinson, Fanselow and Todd (2012) looked at the best strategies for finding a parking space. Moussaid, Helbing, and Theraulaz (2011) considered how simple heuristics guide pedestrian behavior. Hills, Jones and Todd (2012) applied animal foraging models to the study of how people search semantic memory. Hills, Goldstone, and Todd (2008) investigated how searching in physical space influences search in abstract space. In related work, Fu and Pirolli (2007) combined models of human semantic memory and foraging theory to explain how people search for information on the web. During the conference, I hope to understand better the connections between what we have learned about how simple heuristics can get by with limited search and what has been learned about spatial search.

References

Fischer, J., Steiner, F., Zucol, F., Berger, C., Martignon, L., Bossart, W., Altwegg, M., & Nadal, D. (2002). Use of simple heuristics to target macrolide prescription in children with community-acquired pneumonia. *Archives of Pediatrics & Adolescent Medicine* 156(10): 1005–1008.

Gigerenzer, G., & Brighton, H. (2009). Homo heuristicus: Why biased minds make better inferences. *Topics in Cognitive Science 1:* 107–143.

Gigerenzer, G., Hertwig, R., & Pachur, T. (Eds.). (2011). Heuristics: The foundations of adaptive behavior. New York: Oxford University Press.

Hills, T., Jones, M. & Todd, P. M. (2012). Optimal foraging in semantic memory. *Psychological Review 119*: 431–440.

Hills, T., Todd, P.M., & Goldstone, R.L. (2008). Search in external and internal spaces: Evidence for generalized cognitive search processes. *Psychological Science* 19(8): 802–808.

Hutchinson, J. M. C., Fanselow, C., & Todd, P. M. (2012). Car parking as a game between simple heuristics. In P. M. Todd, G. Gigerenzer, & the ABC Research Group, *Ecological rationality: Intelligence in the world* (pp. 454–484). New York: Oxford University Press.

Katsikopoulos, K. V., Schooler, L. J., and Hertwig, R. (2010). The robust beauty of ordinary information, *Psychological Review* 117: 1259–1266.

Luan, S., Schooler, L.J., & Gigerenzer, G. (2011). A signal detection analysis of fast-and-frugal trees. *Psychological Review 118:* 316–338.

Luan, S., Schooler, L. J., & Gigerenzer, G. (2014). From Perception to Preference and on to Inference: An Approach-Avoidance Analysis of Thresholds. *Psychological Review*.

Moussaid M, Helbing D, Theraulaz G (2011) How simple rules determine pedestrian behavior and crowd disasters. *Proceeding of the National Academy of Science USA* 108: 6884–6888.

Fu, Wai-Tat; Pirolli, Peter (2007), SNIF-ACT: a cognitive model of user navigation on the world wide web, *Human-Computer Interaction* 22: 335–412.

From Spatial to Cognitive Search in Humans

PETER M. TODD
Provost Professor of Cognitive Science, Psychology, and Informatics
Indiana University, Bloomington
Email: peter.m.todd@gmail.com

Humans and other animals have specialized cognitive mechanisms adapted for searching for resources in the spatial environment. One commonly-sought external resource is food: When seeking food that comes in spatially-localized patches, such as berries ripening on widely-separated bushes, people (and other species) must decide when the current resource patch (e.g., berry-bush) has been mostly exhausted, and when it is time to leave that patch and move on to a new one. The cognitive and neuropsychological mechanisms underlying this important adaptive spatial search behavior appear to be conserved across many species, from invertebrates to humans (Hills, 2006), relying on dopamine-influenced *area-restricted search* (ARS) algorithms. Recently, there have been suggestions that, in humans at least, similar types of search algorithms may also be used in internal, cognitive searches, such as finding solutions to puzzles or words in memory (Todd, Hills, and Robbins, 2012). If this is so, we may now begin to apply a powerful theoretical tool, foraging theory, to the understanding of human cognition and memory, to the study of human search in online information environments, and even to consideration of pathologies related to dopamine imbalance in search behavior, such as obsessive-compulsive disorder (OCD).

The optimal strategy for deciding when to leave a resource patch during foraging is given by the *marginal value theorem* (Charnov, 1976): Leave a patch when the instantaneous rate of resource return falls below the long-term return rate in the whole environment. Behavioral ecologists have proposed and investigated the performance of sub-optimal, but computationally simpler, rules of thumb such as giving up and leaving a patch after a fixed amount of time (Bell, 1991). Which of these rules perform well depends on the structure of the environment, for instance whether patches are all similar in quality or a few are good and others are poor.

Previous research has found that humans forage in a seemingly optimal fashion for food (Smith and Winterhalder, 1992) or for information in libraries or on the internet (Pirolli and Card, 1999), but has not identified the kinds of simple rules or heuristics being used. We have studied whether humans apply rules similar to those that animals use for deciding when to give up on one resource patch and move to another. One goal was to investigate whether humans may use the same kinds of patch-leaving heuristics both in external resource foraging tasks such as searching for food and in internal cognitive foraging tasks such as searching for words in memory. To test this, we devised two comparable search tasks: first, a (simulated) external search, in which people had to try to catch fish in a pond until they decided that most of the fish were caught and it was time to move on to the next pond (Hutchinson, Wilke, and Todd, 2008); and second, an internal search, in which people had to come up with words that could be formed from a scrambled set of letters (e.g., "LOAF" can be formed from letters in "WLOKFA"), until they decided that they could not think of any more words and it was time to switch to the next set of letters and begin again (Wilke, Todd, and Hutchinson, 2009). We found that people exhibited the same kinds of patch-leaving behavior in both internal and external searches, treating both kinds of environments as though they were patchy and highly variable in quality. This behavior was generated by simple rules implementing forms of area-restricted search. More recently, we have developed new simple computerized games to assess how individuals manage the trade-off between exploiting current resource patches and exploring for new ones (Sang, Todd, and Goldstone, 2011).

Area-restricted search is a common animal foraging behavior involving concentrated local search where resources have been found in the past, followed by a transition to more global search when resources locally become depleted. In an ecological context, animals using area-restricted search will turn more frequently following an encounter with food, temporarily restricting their search around the area where food was just found. In the experimental tasks described above, finding resources also led people to stay longer in a patch. Across species, this transition from exploitative to exploratory search behavior involves brain mechanisms relying on the neurotransmitter dopamine. Similar dopaminergic mechanisms operating between the frontal cortex and striatum in the human brain control goal-directed attention, and are also involved in pathologies of attentional control. This and other evidence supports an evolutionary relationship between animal foraging behavior and goal-directed cognition (Hills, 2006).

In recent work we have shown a direct behavioral relationship between area-restricted search in space and cognition in humans. Using a spatial foraging task followed by a word search task (like that used by Wilke et al., 2009), we found that individuals who were made to forage in patchy spatial distributions (versus diffuse distributions) stayed longer in "word patches" before moving on to the next "patch." We also found that individuals who explored the spatial distribution less widely were also likely to explore less in word patches (Hills, Todd, and Goldstone, 2008). In both cases, perseveration in space was followed by perseveration in cognition. Whether or not this was caused by the former task altering dopamine levels which in turn affected performance in the latter task remains to be tested.

We next explored whether humans search in memory using dynamic local-to-global search strategies similar to those that animals use to forage between patches in space, in ways that correspond to optimal foraging strategies seen for spatial foraging (Hills, Jones, and Todd, 2012). We used the verbal fluency task in which we asked participants to recover from memory as many types of animals as they could in three minutes. People typically produce a burst of items related in some way (e.g. farm animals), then pause, and then produce another burst of related items (e.g. pets). These item clusters represent patches in semantic memory. We modeled memory search over a representation of the semantic search space generated from the BEAGLE memory model (Jones and Mewhort, 2007), using a search process similar to models of associative memory. We found evidence for local structure (i.e., patches) in memory search; furthermore, patch depletion preceded individual's switches between patches using dynamic local-to-global search strategy transitions. The timing of these patch switches was consistent with optimal search policies in space specified by the marginal value theorem, and participants who were more consistent with this optimal policy recalled more items.

All of these results argue for a continuity between the mechanisms that animals including humans use to successfully search for resources in space, and mechanisms that humans use to forage for information, whether externally on the Web, or internally in memory. Ongoing research in our lab is elucidating the exploration and patch-leaving decision mechanisms that people employ when searching through environmental structures created by others, such as the space of "tagged" music or webpages that people build and explore using software tools such as last.fm (for music searching/listening) or Delicious (for webpages). Understanding the spatial origins of these search behaviors can help us design better systems for people to navigate the new forms of resources that we seek.

References:

Bell, W. J. (1991). *Searching behaviour.* London: Chapman and Hall.

Charnov, E.L. 1976. Optimal foraging: The marginal value theorem. *Theoretical Pop. Biology 9*: 129–136.

Hills, T. (2006) Animal foraging and the evolution of goal-directed behavior. *Cognitive Science 30*: 3–41.

Hills, T.T., Jones, M.N., and Todd, P.M. (2012). Optimal foraging in semantic memory. *Psychological Review 119*(2) : 431–440.

Hills, T.T., Todd, P.M., and Goldstone, R.L. (2008). Search in external and internal spaces: Evidence for generalized cognitive search processes. *Psychological Science 19*(8)*:* 802–808*.*

Hutchinson, J., Wilke, A., and Todd, P.M. (2008). Patch leaving in humans: Can a generalist adapt its rules to dispersal of items across patches? *Animal Behaviour 75*(4): 1331–1349.

Jones, M.N., and Mewhort, D.J.K. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological Review 104*: 1–37.

Pirolli, P. and Card, S.K. (1999). Information foraging. *Psychological Review 106*: 643–675.

Sang, K., Todd, P.M., and Goldstone, R.L. (2011). Learning near-optimal search in a minimal explore/exploit task. In *Proceedings of the 33rd Annual Conference of the Cognitive Science Society* (pp. 2800–2805). Boston, MA: Cognitive Science Society.

Smith, E.A. and Winterhalder, B., Eds. (1992). *Evolutionary Ecology and Human Behavior.* New York: Aldine de Gruyter.

Todd, P.M., Hills, T.T., and Robbins, T.W. (Eds.) (2012). *Cognitive search: Evolution, algorithms, and the brain.* Cambridge, MA: MIT Press.

Wilke, A., Todd, P.M., and Hutchinson, J.M.C. (2009). Fishing for the right words: Decision rules for human foraging behavior in external and internal search tasks. *Cognitive Science 33:* 497–529.

Varieties of Spatial Search

BARBARA TVERSKY
Professor of Psychology, Columbia Teachers College
Professor of Psychology Emerita, Stanford University
Email: btversky@stanford.edu

Searching in space occurs frequently and at many levels. Readers might search for the location of this photo. There is the daily search for keys, phones, and glasses; there is search for a coffee shop in the neighborhood, a street sign in the dark, an acquaintance at a party, a seat at the movies, an unseen cry on a playground, the perfect dress, where to insert a screw inside an engine, the source of an unusual smell, the place in a book that discusses a specific idea, information on the web; there is the search for a good design in a vague sketch, for a meaningful pattern in messy data. At a minimum, search in space entails finding a target in a context. Both target and context can vary widely, in richness, in familiarity, in clarity, in discriminability, in specificity, in modality, and more. Search can be guided by habits, by strategies, both relevant and irrelevant (the proverbial search under the streetlight), by cues, perceptual, linguistic, or embodied, real or virtual. Search in space necessitates an interaction between the external, what is perceived in the world, and the internal, how the external is construed, and how the search is conducted. To understand the full range of spatial search is to draw on the full range of perception and cognition.

Our lab has studied many aspects of spatial search across a variety of contexts. Several projects have investigated the qualities of the language used to describe environments, to direct search for a target, or to accomplish a complex task in space. We found that people mix allocentric and egocentric spatial reference systems and take into account the relative cognitive loads of the director and the actor. Other projects have been investigating the roles of gesture, notably pointing, as well as the integrated interactions between gestures and language. We are finding that gestures are preferred to language and disambiguate spatial locations and guide spatial planning and learning more efficiently and successfully than spatial language. Other projects have been investigating qualities of perceptual cues meant to guide search. One project analyzes cues in graphic narratives. Another has been investigating the use of augmented reality (AR) to guide users in maintenance and repair tasks and in navigation. For maintenance and repair, we found that augmentations that seamlessly guided the position and orientation of the body, the head, and the eyes were more successful than a sequence of gauges. Ongoing work is developing interactive AR tools for navigation in space. Still other projects have been studying how different data displays affect sense-making, how designers use ambiguous sketches to find good design solutions, and how various kinds of cues facilitate that process. Especially in these cases, the ambiguity of the external and the intensity of the internal make for a rich paradigm for understanding the conversation between external and internal in spatial search.

Gazetteers and the Library Catalog: Infrastructure for Spatial Search

KATHY WEIMER
Professor, Texas A&M Libraries
Manager, Map & GIS Library
Texas A&M University
Email: k-weimer@library.tamu.edu

The Role of the Gazetteer and Library Records in Spatial Searches

The fields of GIScience and Library/Information Science have numerous overlapping interests, including spatial search. Each discipline is devoted to meaningful description and organization of place names; resulting in gazetteer development and related library catalog data that informs spatial search mechanisms. As with gazetteer development, a robust infrastructure of cataloging standards and codes, built over many decades, has resulted in an extensive place name index, referred to as "geographic name authority records," used as the formal official name in the catalog system. These name authority records function to re-direct searches which may include a variant name to those with the official form of the name and thus, when associated with geographic coordinates has the dual role of a tool and a dataset, as is the role of gazetteers. The larger concepts of place name identification and their conceptual structuring in the library catalog system and their use in spatial searches can and should extend from the library catalog system into GIS, digital gazetteer development and spatial search queries. Now that library catalog data is available through linked open data, discussions should pursue crosswalks or interoperability between catalogs and GIS.

As manager of a map and GIS library, I am keenly aware of humanists' use of digital gazetteers, sometimes enriched by their own data, that structure for interactive scholarly works. However, these scholars seem unaware of the rich resource provided by library catalog records for place names, and by creating their own gazetteers, face many of the same challenges that libraries have sorted through and routinized over many decades. For example, how are regions noted, names disambiguated and changes over time described? How are time periods and events defined and coded? The issues have been discussed and resolved to a great extent in the library community, however, humanities scholars are largely unaware.

The library catalog and place name authority record structure provide a number of possibilities for expanding the utility of spatial search in numerous domains. In approaching the library's geographic name authority records as a digital gazetteer, I suggest that the geospatial community can exploit the vast information contained therein to build or enhance existing digital gazetteers and related spatial search engines. Further, the taxonomies, classifications and standards already existing in library science should be fully utilized by GI-related science and systems to expand and promote spatial search. Crosswalks can be created between geonames, library catalog, and other digital gazetteers, such as the ADL and Getty Thesaurus. Obscure, localized places identified only in a library catalog can be brought into a GIS and national gazetteers.

The Role of Linked Open Data

The library records are available through linked open data formats so can be ingested for various uses. Codes, vocabularies and standards used by Library of Congress have recently been released, found at http://id.loc.gov and through VIAF, the Virtual International Authority File. At the specialist meeting, I would convey the subtleties of the Library of Congress rules for creating official place names and their disambiguation, as well as coding standards, to support those in GIS who may pursue enhanced spatial searches using the LC records for place names.
Much can be gleaned from this existing dataset both in content and organizational methods to enhance larger spatial search initiatives.

Mapping Applications to Search Library, Museum and Archive Collections

In a completely different approach, without regard to library catalogs, mapping applications can provide spatial awareness and spatial search for library collections. I am project director for "Mapping Historic Aggieland" for which georeferenced historic maps of campus and geotagged photos are accessed through a geoserver. Dates of building construction are added, which combined with the maps, form a spatio-temporal browse. Providing a spatial or map-based search to a library, museum or archival collection, will surely become a more popular approach and will bring attention to the geographic aspects of a collection.

Collaborations

In2013, perceiving a need for greater communication and collaboration among those working in the geospatial humanities, I co-founded an international special interest group, "GeoHumanities," within the Alliance of Digital Humanities Organizations (see http://geohumanities.org/). Our inaugural meeting was held this summer on the topic of gazetteers, focusing on their use and development among humanists. Also, I participated recently at a meeting to build a World Historical Gazetteer, where gazetteer interoperability, textual and geometric spatial searches were major discussion topics. From my work with members of the GeoHumanities SIG and the World Historic Gazetteer meeting, I am carrying the humanists' concerns into discussions with librarians across the country to bring awareness about the need for enhancing the geographic name authority records, particularly regarding temporal aspects and geometries, sharing through linked open data in order to expand the use of library catalog records and digital gazetteers for spatial searches. The specialist meeting provides me the opportunity to speak to geographers and others interested in spatial search, about the role, concerns and impacts that libraries may play in this area.

Spatial Search Based on Similarity

MAY YUAN
Geospatial Information Sciences
University of Texas, Dallas
Email: myuan@utdallas.edu

This position paper explores the idea of spatial search based on similarity. People commonly look for places of similar characteristics or offering similar experiences. When one moves to a new location, the person may look for a similar neighborhood to the previous residence. When we observe a phenomenon of interest at a place, we may look for similar places to examine if the phenomenon also exists. There are ample examples that we seek similar places or places providing similar experiences. Where are cities similar to New York? Where can I find a similar event? Where also experience record breaking heat after August in the last 10 years? Where will be likely to become the next Silicon Valley? For the general public, spatial similarity search may suggest candidates for the next travel destination, opportunities to participate in activities or events of interest, and options of places to enjoy food or excitements based on one's fondness of previous experiences. For scientists, spatial similarity search facilitates finding alternative studied areas, identifying places suitable for quasi-experimental designs of geographic factors and consequences, and assessing applicability of models developed from one place to another place.

Conventionally, spatial search on the web seeks specified information within a spatial extent, which could be a proximity around a focal location (e.g., 100 feet from my current location) or a predefined geographic region (e.g., the State of Texas). Many GIS packages are able to support advanced queries that are based on more complicated spatial relationships, such as adjacency and partial containment. Location-based services enable the spatial search for Italian restaurants nearby my current location, or for the origin and destination and how to get there. Moreover, the current web services also allow the National Weather Service to associate weather observations and forecasts with cities to support spatial search of weather measurements and forecast future weather. Nevertheless, spatial search currently depends very much on geographic coordinates to compute distance and spatial relationship. While transportation networks are considered in calculating Manhattan distance for routing on web map services, Euclidean distance is mostly common for spatial search in proximity.

As Tobler's law suggests that similar things are closer in space, spatial proximity could imply spatial similarity. As such, distance can serve as one factor in assessing similarity. Perhaps, however, of the most interesting are places afar but exhibit a high degree of similarity. How an indigenous village in South East Asia may have similar cultural practices and spiritual artifacts to what takes place in a Maya tribe in South America may be of great interest to anthropologists or eco-tourists. How City A in Arizona facing water shortage may find City B of similar characteristics in Texas and learn from City B its successful experiences in mitigating water issues.

Realization of spatial similarity search is challenging. However, there may be three areas worthy of research exploration:

1. Mash-up and Spatial Analytics: Mash-up plays an important role in geospatial services and web GIS. It helps build web GIS applications by fusing web map services, web feature services, and web geoprocessing services to make interactive spatial search more meaningful and relevant through contextualizing and spatially relating geographic information. Rapid developments in Web GIS, Cyber GIS, and Cloud GIS have pushed more and more spatial tools available online for attribute query and spatial analysis; most of these tools mimic what are available in a desktop GIS. A bold thinking is enabling spatial search that explores relevant geospatial services and retrieve, compute, and summarize outcomes automatically. A high school student who visits a college and likes the college town may click on the college town and ask where similar college towns are in the U.S. Currently, we can build geoprocessing models on the web to retrieve the college town attribute data based on which to retrieve other college towns of similar population, and median income, etc. The geoprocessing model may include other data, such as shopping centers or bike paths for analysis, but the model only runs on the specified web services in the application. What is envisioned here is that spatial search expands to the mash-up phase so that it is part of the search algorithm to find relevant geospatial services input to the geoprocessing model. Current spatial search for relevant geospatial services is normally done by the names, keywords, and some meta information about the services. What is considered here is the ability to spatial search for relevant features in a geospatial service before determining which web service is relevant.
2. Spatial search with text analytics: Many places already have much information on the web. City government web sites have basic city descriptions about population, geography, history, landmarks and major economic activities. Many cities also have entries in Wikipedia. Some web sites include questions and responses about places. Much information about places (especially historical places, tourist sites, and places with major sports events) is also available from travel pages, online tourist guides, and news media. Many magazines (like *Economist*) and non-profit organizations (like livability.com) rank cities with detailed information. These are good sources for text analytics to identify city characteristics that are useful to support searching for places with desired properties or comparing places based on specific parameters. A bold vision here is a live digital gazetteer that is continuously being updated, enriched, and indexed by text analytics agents with sourced materials from the web resources or through crowd sourcing. Toponym resolution can pose additional challenges in identifying place names, disambiguation of place names with other named entities, and determining coordinates for identified places. Local place names vs official place names shall also be considered in such a gazetteer.
3. Geographic genomics: What are fundamental and extended parameters characterized a place? Urban geographers differentiates the concepts of "site" and "situation in the study of settlement patterns and what factors lead a settlement to becoming a large city over time or remaining as a small village. Factors of sites and situations may be conceptualized as "geographic genomics" and, as such, similarity among places can be assessed through linear alignments of their geographic genome sequences. In general, site factors include landforms, climate, vegetation types, soil quality, wildlife, water and other natural resources. Site factors determine the innate capacity and affordance of the site. Situation, on the other hand,

> concerns with the location of a place relative to its surroundings and other places. Common situation factors include the extent of a place's connections with other places, accessibility to resources from outside, and the place's political, cultural, and economic roles now and in the history. Can we use these site and situation factors to construct geographic genomes? Should we take a different approach to identify genes of a place to support search for similar places? Are there basic sets of geographic genes that can partially or as a whole to help determine the degree to which two places are similar and furthermore to determine if a model developed from one place can be transferable to another place? Or to what degree we can trust the outcomes from such a model transfer? Moreover, if we can find two places with identical geographic genes but one or two, can we apply quasi- experimental design to determine causal relationships of the one gene on different outcomes resulted from the same stimulus (e.g., one community successfully rebounded from a F4 tornado, but another community did not).

This position paper considers spatial search a fundamental research that quests for the essential characteristics of places and based on the essential characteristics we can assess spatial similarity of places, design quasi-experiments to reveal factors that contribute to different outcomes, and determine the degree to which a model developed in one place can be transferable to another. There are vast amounts of resources available on the web. Spatial search research should go beyond coordinates and spatial proximity to better utilize web resources for geographic knowledge production.

SPECIALIST MEETING

Spatial Search

December 8–9, 2014

Upham Hotel, Santa Barbara, California

Information search has become a fundamental element across the spectrum of human activities. Search engines process billions of queries each day and influence the visibility and accessibility of online content. Scientists search for meaningful patterns in massive datasets, while consumers search for products and services in a growing pool of options.

The Center for Spatial Studies at the University of California, Santa Barbara is hosting a specialist meeting, entitled *Spatial Search*. The meeting will bring together academic and industry representatives from **computational, geospatial,** and **cognitive** disciplines with interest in focused discussions to illuminate these matters and contribute toward the development of an interdisciplinary research agenda to advance spatial search, both from scientific and engineering viewpoints.

Participants will address such questions as:

- What are the current challenges in spatial search?
- What are the limits of spatial indexing?
- What kinds of spatial search are utilized in geo-spatial domains?
- What disciplines' functionality is missing in current Geographic Information Systems?
- What do we know about how humans conceptualize and perform information searches?
- How do human perform search of memory and of visual and aural stimuli?
- Can hypotheses and insights from the cognitive and neurosciences inform computational and geospatial search techniques?

For more information, see http://spatial.ucsb.edu/2014/spatial-search

Motivation

Operating at two levels, there is a spatial component at the core of search: **(1)** Search technologies rely on a spatial metaphor; we talk about going to our favorite websites to help search for small fragments in an overwhelmingly large space of documents, images, and videos. **(2)** Geographic space is increasingly important to index information and refine search strategies, relying on the geo-location of entities to assess their relevance. While the spatial dimension of search is pervasive and foundational to many disciplines, it has never been adequately analyzed.

Speakers:

Stuart Card
Dept. of Computer Science
STANFORD UNIVERSITY

Miguel Eckstein
Dept. of Psychological & Brain Sciences
UNIVERSITY OF CALIFORNIA, SANTA BARBARA

Vanessa Murdock
MICROSOFT

Ross Purves
Dept. of Geography
UNIVERSITY OF ZURICH

Subject to approval, limited funding for travel and accommodation costs will be available to invited participants.

Goals

The overarching goal of this specialist meeting is to put spatial search at the forefront of discussion, tackling research questions from three complementary strands: computational, geospatial, and cognitive strands.

Means

About 35 experts from academia and industry will convene to share and develop visions, insights, and best practices. Plenary presentations and intense exchanges in small breakout discussion groups offer opportunities for knowledge transfer.

Call for Applications

To apply, please submit a biograph form résumé with a photograph and a short position paper (in one Word document) discussing your perspective on the subject by **September 20, 2014**. Participants will be selected by the organizing committee and notified of their acceptance by **October 10**.

Please go to http://spatial.ucsb.edu/2014/spatial-search/submissions to submit an online application.

ORGANIZERS:
Werner Kuhn, Mary Hegarty, Andrea Ballatore
(Center for Spatial Studies, University of California, Santa Barbara)
Ed Parsons (Google)

spatial@ucsb
CENTER FOR SPATIAL STUDIES
Google