Applying Space Syntax to Online Mapping Tools

Yandi Li
Universitat Politecnica de Catalunya
liyandi007@163.com

Nicola Barbieri
Tumblr
barbieri@yahoo-inc.com

Daniele Quercia
Bell Labs, Cambridge, UK
quercia@cantab.net

ABSTRACT

To walk around the city, individuals use mobile mapping services, and such services mostly suggest shortest routes. To go beyond recommending such walkable routes, we propose a new framework for automatic wayfinding for pedestrians. This framework tackles two main drawbacks from which past work suffers, namely coarse-grained representation of space and absence of contextual dynamics. We model the human tendency to regularize space by borrowing a spatial representation, Space Syntax, from the discipline of Architecture. Moreover, the proposed framework accounts for contextual dynamics of individual streets by predicting the popularity of each street under different contexts (e.g., at a given time, with a certain weather condition). Using Foursquare check-ins (i.e., whereabouts of the users of the popular location-based service) and publicly available weather data, we validate our framework in the entire city of Barcelona. We find that, with paths slightly longer than the shortest ones, our framework is able to accommodate our mental topography and effectively capture contextual changes.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous

Keywords

Experimental Study, Path Recommendation, Urban Informatics

1. INTRODUCTION

More than half the world's people now live in cities and, by 2050, more than 70% will do so. At a time when millions of individuals are moving to cities, understanding the pragmatics of designing navigation tools is a major consideration.

Many behavioral scientists are aware of the principles that underlie successful wayfinding [3, 11, 18], and those principles have played some role in the design of navigation tools but usually a secondary and accidental one. They have often taken a back seat to other concerns, which reflect an over-reliance on efficiency. For example, conventional wisdom holds that, since time is a commodity in short supply, navigation tools have to limit themselves to recommending the most efficient (shortest) routes. For simplicity's sake, such tools should also shy away from the complexity of cities. We all know that activities in the city shift cyclically and progressively as contextual factors (e.g., time of the day, weather) change [17], and yet wayfinding tools have little awareness of city dynamics.

A case in point is the difference between day and night. Take the venues Foursquare users in Barcelona go to at night and those they go to during the day. By connecting those venues, one can readily observe, for the very two same points of 'Arc de Triomf' and 'Estacio de Franca', two very distinct paths (Figure 1).

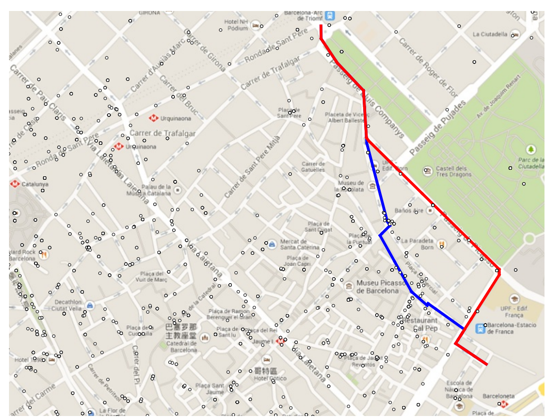

Figure 1: Same destination, different beaten paths for day (red) and night (blue). At daylight, Foursquare users in Barcelona go through the park. At night, through the little streets full of cocktail bars.

Our goal is to propose a new way of navigating the city that is aware of how urban rhythms change with contextual factors. The main research contribution consists of a proposal that incorporates principles of successful wayfinding and accommodates contextual changes just by adding a few extra minutes to walking time. More specifically, we make four main contributions:

- We model contextual dynamics by adapting state-of-the-art techniques such as Factorization Machines and ranking optimization techniques to our problem (Section 3.1). The idea is to model how the interestingness of different points in the city (e.g., Foursquare venues) changes with contextual variables such as time of the day or weather.
- We model space by incorporating principles of successful wayfinding. We do so by borrowing a modeling technique from the field of Architecture called space syntax (Section 3.2). This technique mirrors our own mental maps by representing

WSDM 2017, February 06-10, 2017, Cambridge, United Kingdom

ACM 978-1-4503-4675-7/17/02…$15.00
DOI: http://dx.doi.org/10.1145/3018661.3018722

space as a series of simple viewpoints (nodes) and connections between them (edges). At the core of our research contribution, this way of modeling space can be integrated with any context modeling other than the one proposed in the previous point.

- We combine the previous two points and propose a new framework that models contextual dynamics within the space syntax representation (Section 3.3).
- We evaluate our framework using Foursquare and weather data in the city of Barcelona (Section 4). We find that our framework is able to suggest contextually relevant routes, it is able to do so without any considerable walking overhead, and it might be used not only in central parts of the city but also in low-density areas.

The main novelty of our work is about embedding context-aware desirability of street segments within the space syntax framework. This is far from being trivial for three main reasons. First, despite route finding has already exploited space syntax [4], no working solution that integrates space syntax with current mapping technologies in the commercial world has been engineered yet. Second, it is unclear how to solve the data sparsity problem (i.e., street desirability scores might not always be available). To automatically fill missing data, we will propose to use matrix factorization to learn dependencies directly from the data (e.g., desirability scores at specific times of day might be related and, as such, that information might be used to fill potential missing desirability scores). Third, we do not know how to fine tune the parameters of a framework that combines space syntax with matrix factorization for route finding.

2. RELATED WORK

Wayfinding has been a popular research topic [3]. Considerable research has gone into how to integrate human psychological aspects [19, 32, 34] and contextual factors into wayfinding tools (by contextual factors, those researchers mainly meant the conditions of the built environment [15, 31]). Some research has also gone into how to leverage psychological aspects to design better ways of visualizing navigation information on mobile devices [20].

Early work in context-aware recommendation has used contextual information either for pre-processing, whereby context is used for data selection, or for post-processing, whereby context is used to filter recommendations [1]. Later work has, instead, focused on contextual models that integrate recommendations, users, and contexts at the same time. The main state-of-the art approaches are two: tensor factorization [14, 30], and factorization machines [27, 29]. Those approaches have been initially designed to work on explicit ratings but have been recently extended to work on implicit feedbacks as well [5]. We capitalize on those recent insights to explore the applicability of factorization machines to our problem.

A lot of research work has gone into the problem of trip recommendation. Its goal is to include good venues in a trip. Past research proposals mainly differ in what they mean by good. For Gionis *et al.*, good venues are those that respect time/distance constraints and are ordered in a reasonable way (e.g., lunch at a restaurant might be followed by a coffee shop) [7]. In this context, a good path is a set of venues that are desirable for a given user or user class (e.g., tourists vs. locals). For Lu *et al.*, good venues are those that meet specific constraints (e.g., time budget, a user's liking, need for diversity) [16]. More recently, for Quercia *et al.*, good venues are places in the city that are quiet or beautiful, or that make people happy [24]. To create paths out of those concepts, they modeled the city by dividing it into walkable cells. Within each cell, perceptual factors such as beauty, quite and happiness were crowdsourced [25]. The authors conceded that their work has two main limitations though: an over-simplified spatial representation based on cells that does not allow for a spatial analysis finer than cell-level; and disregard for the effect of contextual factors.

To sum up, for simplicity's sake, past work on route recommendation and wayfinding has relied on coarse-grained spatial representations and on over-simplified modeling of contextual dynamics.

Table 1: Main Symbols.

Symbol	Description
$\mathcal{V}$	set of venues ($\lvert\mathcal{V}\rvert = N$);
T	number of venue categories;
$t(v)$	category for venue v;
$sg = (s, t)$	axial segment from location s to t;
$Ax(sg_f)$	axial line for the segment sg_f;
$len(sg_f)$	lenght of the segment sg_f
$G = (S, E)$	segment graph defined over segments S, $e_{f,g} \in E$ means that segment f and g are adjacent;
$P_{s,t}$	path (sequence of segments) from location s to t;
$\mathbf{c}$	context descriptor (binary M-dimensional array);
C	set of observed contexts;
$b(sg_f; \mathbf{c})$	contextual interestingness of segment sg_f;
$\hat{y}(v; \mathbf{c})$	predicted # check-ins on venue v in context $\mathbf{c}$;
$w_{f,g}(\mathbf{c})$	contextual cost associated to $e_{f,g}$;
$d_{f,g}$	metric distance from sg_f to sg_g;
$\theta_{f,g}$	angual change between sg_f to sg_g;
$b(sg_f; \mathbf{c})$	interestingness of segment sg_f under context $\mathbf{c}$.

3. RECOMMENDATION FRAMEWORK

To move the research forward, we propose a framework for contextual route recommendations within a principled spatial representation. A framework consists of algorithms and models, each of which can be replaced. In that way, a framework is more general than its composing models/algorithms.

For convenience, Table 1 collates the main symbols we will use in this section.

Problem Definition. Our goal is to determine, given starting and ending locations (s, t), a path $P_{s,t}$ that is short and interesting under the specific context, where by context we mean, time of the day, time of the week, and weather conditions. To this end, our framework (a) models how each path's dynamics change with context; (b) embeds each path in a spatial representation mirroring human mental topography; and (c) recommends a path by combining contextual dynamics with the specific representation of space.

3.1 Model of Contextual Dynamics

Our smallest unit of analysis is that of individual venues in the city. Let $\mathcal{V} = \{v_1, \cdots, v_N\}$ denote a set of N venues, where each venue v_i is a physical location (e.g., business, residence), which can be categorized into T types (e.g., food, entertainment, shop), and $t(v_i)$ denotes its category.

We are interested in modeling how a venue's interestingness changes with contextual factors. Each context can be represented as a m-dimensional binary array $\mathbf{c} = \{c_1, \cdots, c_M\}$; for instance, considering the following contextual dimensions {day, night, weekday, weekend, rainy, sunny}, the contextual descriptor $\mathbf{c} = (1, 0, 0, 1, 0, 1)$ corresponds to "day, weekend, sunny". Venue iv_i's interestingness in context $\mathbf{c}$ needs to be predicted. That is because in reality, we are not likely to observe the interestingness of all venues under all contexts and, as such, most of the information will be missing.

Thus, upon observed interestingness information, we predict the missing one. More formally, we predict $\hat{y}(v_i; \mathbf{c})$, which is the estimated interestingness of venue v_i under context $\mathbf{c}$. To make such predictions, we use the Factorization Machine framework. We chose it because it is an elegant formalization that has been shown to work extremely well for numerical predictions [27]. The parameter estimation phase is based on learning algorithm which, by exploiting pairwise comparisons, is directly optimized for ranking. The combination of latent dimension modeling and ranking-based learning allows the model to learn the popularity of venue-context pairs that are *unseen* at training phase. The underlying assumption behind this framework is that the predicted response for each pair $(v_i, \mathbf{c})$ can be decomposed as a linear combination of the interactions between input variables and K latent factors. In our case, we consider the following model:

$$\hat{y}(v_i;\mathbf{c}) = \mathbf{w}^0_{v_i} + \mathbf{w}^T_{t(v_i)} + \sum_{j=1}^{M} \mathbf{c}_j \cdot \sum_{k=1}^{K} \left(\mathbf{U}^V_{v_i,k} + \mathbf{U}^T_{t(v_i),k}\right) \cdot \mathbf{U}^C_{j,k},$$

where:

- $\mathbf{w}^0 \in \mathbb{R}^{1\times N}$ is the estimated interestingness for each venue without considering context and venue type;
- $\mathbf{w}^T \in \mathbb{R}^{1\times T}$ is the estimated interestingness for each venue type without considering context;
- $\mathbf{U}^V \in \mathbb{R}^{N\times K}$ reflects all venues' positions in the latent space;
- $\mathbf{U}^T \in \mathbb{R}^{T\times K}$ reflects all venue types' positions in the latent space;
- $\mathbf{U}^C \in \mathbb{R}^{M\times K}$ reflects all contexts' positions in the latent space.

Given that $\Theta = \{\mathbf{w}^0, \mathbf{w}^T, \mathbf{U}^V, \mathbf{U}^T, \mathbf{U}^C\}$ is the set of model parameters, and $\mathbf{Y}$ is the set of observed interestingness values, we optimize Θ according to the following loss:

$$\Theta^* = \arg\max_{\Theta} \sum_{\substack{(v_i,\mathbf{c}),(v_h,\mathbf{c})\in\mathbf{Y}\\ y(v_i;\mathbf{c})>y(v_h;\mathbf{c})}} \ln\sigma(x_{i,h,\mathbf{c}}) - \lambda \left\|\Theta\right\|^2,$$

where $x_{i,h,\mathbf{c}} = \hat{y}(v_i;\mathbf{c}) - \hat{y}(v_h;\mathbf{c})$, $\sigma(x) = \frac{1}{1+e^{-x}}$ is the logistic link function, $\|\Theta\|$ is the L2-norm of parameters of the factorization model. As a regularizer, we assume a prior distribution $\Theta \sim N(0, \lambda I)$, where λ is a regularization parameter (we experimentally found the best value of λ to be $\lambda_{w^0} = \lambda_{w^T} = 0.1$, $\lambda_{U^C} = \lambda_{U^T} = 0.01$, $\lambda_{U^V} = 0.0025$).

When making predictions, this schema (which, as a matter of fact, can be considered as a contextualized version of *Bayesian Personalized Ranking* (BPR) [28]) does not focus on minimizing the error loss to score individual venues (as prediction algorithms usually do) but ensures that venue pairs are correctly ranked. To learn the parameters of the model we apply a stochastic gradient ascent procedure, which consists of multiple iterations. At each iteration, it selects a venue v_i and corresponding training sample $(v_i, \mathbf{c})$, samples another venue v_h at random (a challenger) such that $y(v_i;\mathbf{c}) > y(v_h;\mathbf{c})$, and updates a parameter Θ_p to be:

$$\Theta_p \qquad \Theta_p + \alpha\left(\frac{e^{-x_{i,h,\mathbf{c}}}}{1+e^{-x_{i,h,\mathbf{c}}}} \cdot \frac{\partial}{\partial\Theta_p} x_{i,h,\mathbf{c}} - \lambda_p \cdot \Theta_p\right).$$

This translates into setting the partial derivative $\frac{\partial}{\partial\Theta_p} x_{i,h,\mathbf{c}}$ to, say, 1, if we are updating the parameter $\mathbf{w}^0_{v_i}$. More broadly, the partial derivative is set to take different values depending on which parameter is to be updated:

$$\begin{cases} 1 & \text{if } \Theta_p = \mathbf{w}^0_{v_i},\ \mathbf{w}^T_{t(v_i)} \\ -1 & \text{if } \Theta_p = \mathbf{w}^0_{v_h},\ \mathbf{w}^T_{t(v_j)} \\ \sum_{j=1}^{M} c_j \cdot \mathbf{U}^C_{j,k} & \text{if } \Theta_p = \mathbf{U}^V_{v_i,k},\ \mathbf{U}^T_{t(v_i),k} \\ -\sum_{j=1}^{M} c_j \cdot \mathbf{U}^C_{j,k} & \text{if } \Theta_p = \mathbf{U}^V_{v_h,k},\ \mathbf{U}^T_{t(v_h),k} \\ \mathbf{U}^V_{v_i,k} - \mathbf{U}^V_{v_h,k} + \mathbf{U}^T_{t(v_i),k} - \mathbf{U}^T_{t(v_h),k} & \text{if } \Theta_p = \mathbf{U}^C_{j,k}\text{, where } c_j = 1 \\ 0 & \text{otherwise.} \end{cases}$$

Those formulae have been derived from multiple research papers that have never been collated together before [28, 27].

3.2 Model of Space

After characterizing the dynamics of venue interestingness, we need to place them in space. To model space, one could resort to a traditional representation inspired by Euler's seven bridges of Konigsberg: the urban layout is modeled as a network whose nodes are intersections, edges are roads, and edge weights reflect cost values. This representation is called *primal* graph and is often used for planning policies of cities.

However, such a representation does not account for the way individuals navigate space. That is why, in the late 1970s, Hiller *et al.* developed the *space syntax* framework [9]. Space syntax aims at providing a graph representation of the urban layout that accounts for the way people deal with space and navigation. It simplifies the spatial geometry by reducing complex spaces into sets of points and lines [3]. The core of this methodology is the *axial map*: each open space (e.g., street, square) is approximated by a minimal set of straight lines (street segments or, in technical parlance, *axial lines*), and their connections reflect the elements that are directly visible by humans. Then, in a corresponding *dual* graph representation, each axial line becomes a node, and each intersection between any pair of axial lines becomes an edge [22].

Instead of using the original axial map representation, we opt for a more recent formulation called *Angular Segment Analysis (ASA)* [35]. This breaks axial lines into segments and records the angle between them. *ASA* is becoming a popular tool in urban planning studies and has been introduced to deal with representational problems the traditional axial lines have [26]. In *ASA*, an axial segment $sg = (s, t)$ is a straight line connecting two locations s and t that are on the same street and are reachable by walk without taking turns. The urban layout is then represented as a *segment graph* $G = (S, E)$ (Figure 2), where S is the set of axial segments, and $E \subseteq S \times S$ specifies the adjacency relationships between segments: $e_{f,g} \in E$, if the segments sg_f and sg_g intersect in one of their endpoints. As a final step, we need to associate a cost with each edge. As it is often done in space syntax, we adopt two definitions of cost. These reflect two distinct notions of physical distance: *metric* distance $d_{f,g}$, and *angular change* $\theta_{f,g}$ (examples of those two quantities are highlighted in yellow in Figure 2). The metric distance between adjacent segments is calculated as half the sum of their length: $d_{f,g} = (len(sg_f) + len(sg_g))/2$, where $len(sg_f)$ is the length of the f-th segment. The angular change, instead, is proportional to the angle of incidence of two segments at the intersection and is normalized in the interval $[0, 1]$ (θ is 0, if there is no turn; 0.5, if there is a 90° turn; and 1, if the turn is 180°). By considering the angular change as a notion of cost, space syntax models the idea that a route with many changes in direction is generally perceived to be longer than it actually is: each turn brings into view a new set of physical elements, which form a new region in people's abstraction of space [3].

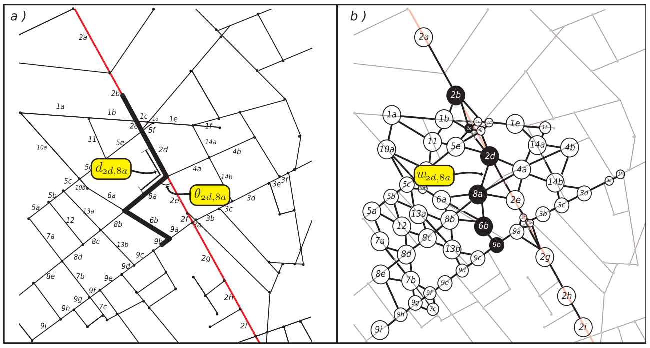

Figure 2: Axial segments of "La Rambla" in Barcelona (left panel) and corresponding segment graph representation (right panel). In both panels, the path from node `2b` to `9d` is in bold. In the left panel, the metric distance d and angular change θ between nodes `2d` and `8a` are further highlighted. In the right panel, instead, the contextual cost w is highlighted.

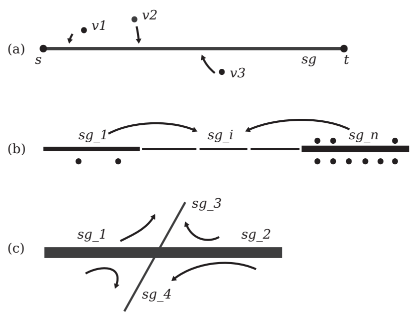

Figure 3: A segment is considered to be interesting if three elements related to it are interesting too: (a) the venues on it; (b) the segments on the same axial line (segments that are visible from it); and *(c)* the segments crossing it.

3.3 Path Generation

Having the segment graph G, for any source location s and a destination t, we should now find a path $P_{s,t} = \left[sg^{(1)} = (s, \cdot), \cdots, sg^{(n)} = (\cdot, t)\right]$. As a path is a sequence of adjacent vertices (segments) from s to t, one can easily find it on G using Dijkstra's algorithm (obtaining the optimal solution) or A* algorithm [8] (settling for a computationally efficient solution). However, either way, G needs to have weights on all its edges, and each weight should reflect the cost of walking from one street segment to the consecutive one in a very specific context. To this end, we associate each edge $e_{f,g} \in E$ with a contextual cost $w_{f,g}(\mathbf{c}) \in \mathbb{R}^+$. This cost quantifies the difficulty of walking from segment sg_f to segment sg_g:

$$w_{f,g}(\mathbf{c}) = f(d_{f,g}, \theta_{f,g}, b(sg_f; \mathbf{c}), b(sg_g; \mathbf{c})).$$

In plain English, the difficulty of walking from sg_f to sg_g depends on the metric distance $d_{f,g}$ one has to walk, on the angular change $\theta_{f,g}$ one experiences, and on how (un)interesting the segments sg_f and sg_g tend to be in context $\mathbf{c}$.

Since quantifying metric distance or angular change is straightforward, only the segment's interestingness remains to be defined. We do so with three iterations (Figure 3).

Step (a) A segment is contextually interesting if the venues on it[1] are contextually interesting:

$$b^0(sg_f; \boldsymbol{c}) = \sum_{v_i \in V(f)} \hat{y}(v_i; \boldsymbol{c}),$$

where $V(f) \subseteq \mathcal{V}$ is the set of venues on segment sg_f.

Step (b) A segment is contextually interesting if "collinear" street segments are contextually interesting:

$$b^1(sg_f; \mathbf{c}) = b^0(sg_f; \mathbf{c}) + \sum_{\substack{sg_g \neq sg_f \\ Ax(sg_f) = Ax(sg_g)}} b^0(sg_g; \mathbf{c}) \cdot \exp(-\frac{d^2_{f,g}}{2\sigma^2}),$$

where $Ax(sg_f)$ is segment sg_f's axial line. The idea is that a segment propagates its influence to other nearby segments which are directly visible by humans [10]. By collinear, we mean segments that are on the same axial line. Since one pair of segments might be closer than another, we need to discount by distance, and that is what σ does. Experimentally, through a grid search, we found the best value of σ be equal to 200 (at that value, the argument of $\exp()$ for a segment that is 235 meters away equals 0.50).

Step (c) Finally, to a lesser extent, a segment is contextually interesting if adjacent street segments (event those not on the same axial line) are contextually interesting:

$$b(sg_f; \mathbf{c}) = b^1(sg_f; \mathbf{c}) + \lambda \left(\sum_{\substack{sg_g \text{ is adj to } sg_f \\ Ax(sg_g) \neq Ax(sg_f)}} b^1(sg_g; \mathbf{c}) \right),$$

where $\lambda \in [0, 1]$ controls the influence of adjacent street segments (λ is set to 0.25 in our experiments). This third step captures the intuition that people may walk on a street simply because there is something interesting "just around the corner".

Having defined all the composing elements of $w_{f,g}(\boldsymbol{c})$, we now need to combine them. We do so based on findings about good wayfinding. Cognitive scientists have found that people tend to minimize changes in direction when walking to destination [11, 18, 21]. However, not everyone minimizes turns. Holscher *et al.* have suggested that, when walking in a neighborhood, residents are comfortable with frequent changes of direction, while visitors mostly rely of visible environmental features and tend to turn as little as possible [12]. Based on this consideration, we define the contextual cost as a combination of metric distance, angular change, and interestingness:

$$\begin{aligned} w_{f,g}(\boldsymbol{c}) \quad = \quad & d_{f,g} \underbrace{(\theta_{f,g}(1 - \alpha_u) + 1)}_{\text{part I}} \cdot \\ & \underbrace{\exp\left(-\frac{(b(sg_f; \boldsymbol{c}) + b(sg_g; \boldsymbol{c}))^2}{2\gamma^2}\right)}_{\text{part II}}. \qquad (1) \end{aligned}$$

To clarify that formula of contextual cost, consider that it consists of two parts.

Part I includes angularity, which is considered be a multiplier on the metric distance. That is because it has been shown that a depth-measure in the segment graph that multiplies angular change with metric length is able to predict traffic flows and pedestrian flows better than what standard axial analysis does [35]. To balance metric distance and angular change, we introduce $\alpha_u \in [0, 1]$, which reflects the extent to which the user is familiar with the city or neighborhood: for first-timers, α_u could be set to 0, and angular change θ would have the highest effect; for long-time residents, α_u

[1]Venue v_i is considered to be on segment sg_f, if v_i's geographically closest segment is sg_f.

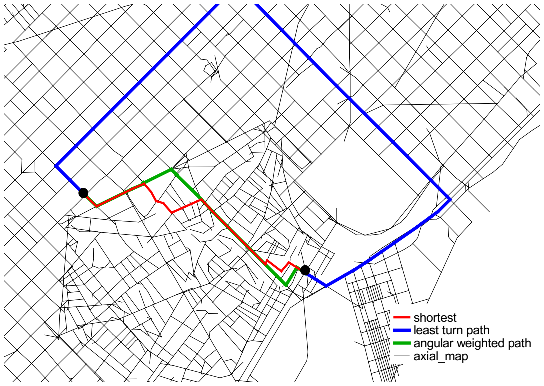

Figure 4: Paths generated for different combinations of metric distance and angular change on the map of Barcelona: the green path combines the two, the red path considers only the metric distance, and the blue one considers only the angular change.

could be set to 1, and angular change would have no effect. To visually illustrate the idea behind α_u, we generate three paths from the same source and destination (Figure 4): the green path is generated by setting $\alpha_u = 0$ (metric distance and angular change are combined), the red by setting it to 1 (only metric distance is considered), and the blue by setting it to 1 while neutralizing the effect pf metric distance (i.e., while setting $d_{f,g}$ to 1). In the green path, the shortest path's (the red path's) zig-zagging is smoothed out by the combination of metric and angular change.

Part II includes the interestingness scores of the two segments and considers them, again, to be multipliers on the metric distance. Intuitively, the overall distance should be discounted if the segments are considered interesting under the given context. This simple design principle can be achieved by combining the two interestingness scores with a Gaussian kernel. The more interesting the segments, the closer the kernel function to 0, and the smaller the perceived distance over the actual one. If the segments are not interesting, the kernel is 1, and the perceived distance corresponds to the actual one. The weight of the multiplier is controlled by a strolling factor γ, which allows for recommended paths that are long (low γ) or short (high γ), as Figure 5 exemplifies (the Figure reflects the results of different runs on real data). For $\gamma \to 0$, a small increase of interestingness results into the generation of longer paths. On the other hand, for $\gamma \to \infty$, the kernel function is close to 1, and that results into the generation of paths very similar to the shortest one. γ can be thus used to enforce constraints about extra-walking time. A user might find desirable to walk a path longer than the shortest one, but only if the extra-walking time does not exceed 10 minutes. To meet this requirement, the strolling coefficient γ has to be properly set. This is done by computing paths for increasing values of γ, and selecting the first one that meets the user's desired extra-walking time.

4. EVALUATION

The goal of our proposal is to generate paths tailored to specific contexts (e.g., for daylight as opposed to for night, for sunny rather than rainy days). To ascertain the effectiveness of our framework at meeting this goal, our evaluation ought to answer four main questions:

- *Recommendation Desirability.* To which extent are our recommendations able to suggest contextually-relevant paths ?
- *Walking Overhead.* How much longer do they take to go much more contextual-relevant paths?

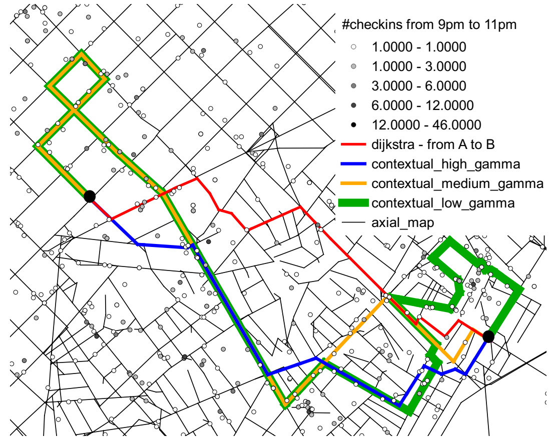

Figure 5: Paths generated for different values of the strolling coefficient γ. As γ increases (i.e., from the green, to the yellow, to the blue paths), stroll is reduced and the resulting path's length is similar to the shortest path's (red path's).

- *Area Density.* Do they work in low-density areas of the city as well?
- *Computational Overheads.* What time and storage overheads does our framework entail?

To see whether our framework effectively suggest recommendations tailored to context, and it does so without resulting into extra walking time, we need to implement and test it in a real city. We choose the city of Barcelona, not least because, for it, we have access to axial maps, contextual data, and changes of street interestingness with time and weather conditions. Next, we will evaluate whether our proposal recommends contextually-relevant paths and, in so doing, how it compares to competitive baselines. We will then ascertain how much longer it takes to find paths that are best in specific contexts, and whether it is effective not only in dense areas (e.g., town center) but also in low-density neighborhoods. Finally, we will test whether our framework's back-end could run in a real-time fashion.

4.1 Experimental Setup

4.1.1 Datasets

To determine whether our route recommendations are tailored to context, we need to have real data on how a venue's interestingness changes with context. To this end, we resort to a Foursquare dataset released by [2]: 22,387,930 Foursquare check-ins collected from September 2010 to January 2011. From these check-ins, we extracted those that happen to be in Barcelona: roughly 60K check-ins in 7364 places from 1690 distinct users. In our experiments, we consider that the number of check-ins at a place is a good proxy for quantifying the place's interestingness, as previous work has shown [33]. However, beyond experiments' sake, we have to stress that our framework is general enough to accommodate notions of popularity other than the presence of check-ins. As a proxy for venue v_i's interestingness in context $\mathbf{c}$, we could consider the venue's beauty score [25], its walkability score [23], or its fit with user interests [16] and needs [7].

The measure of interestingness has then to be studied as context changes. By context, in our experiments, we mean two factors: time and weather. As for time, we consider three temporal granularities: hour, day of the week, and month. We discretize each week into seven days (from Monday to Sunday), the year into 12 (months), and the day into 10 (hour slots), according to a *equal-frequency* binning principle. The 1^{st} hour slot goes from midnight to 7am, the 2^{nd} from 7am to 9am, the 3^{rd} from 9 to 11, 4^{th} the hour starting at 11am, 5^{th} from 12am to 2pm, 6^{th} 2pm to 4pm, 7^{th} 4pm to 6pm, 8^{th} the hour starting at 6pm, 9^{th} from 7pm to 9pm, and 10^{th} from 9pm to 11pm.

The second contextual feature for which we collect data concerns weather. We collect weather data from the Weather Underground[2] and group similar weather conditions together based on quantile frequency. In so doing, we identify 10 main weather conditions[3]. We then stratify check-ins according to weather conditions and time of the day.

4.1.2 Metrics

To measure the extent to which a path recommended by our proposal is desirable, we need to have a metric that reflects path desirability in the first place. A natural but primitive metric is to sum the interestingness scores of all the venues on the recommended path. However, the problem with this naive approach is that the longer the path is, the greater the sum is likely to be. Thus, under such metric, recommending extremely long paths would be desirable. To fix that, we normalize the desirability of path P by its length and call this quantity *raw desirability*:

$$\textit{raw desirability}(P;\boldsymbol{c}) = \frac{\sum_{v \in N(P)} y(v;\boldsymbol{c})}{\sum_{s \in P} len(s)}, \tag{2}$$

where $N(P)$ is the set of venues on (or in the vicinity of) path P, $y(v;\boldsymbol{c})$ is venue v's interestingness (e.g., number of check-ins) for a given context c, and $\text{len}(s)$ is the metric length of street segment s.

Of course, that definition of desirability is not the only feasible one. Variants are possible but ideally should not greatly affect the results. To ensure that, we run our experiments with three variants as well:

- $\textit{raw desirability}^1(P;\boldsymbol{c}) = \frac{avg_{v \in N(P)}\{y(v;\boldsymbol{c})\}}{\sum_{s \in P} \text{len}(s)}$
- $\textit{raw desirability}^2(P;\boldsymbol{c}) = \frac{\sum_{s \in P} \max_{v \in s} y(v;\boldsymbol{c})}{\sum_{s \in P} \text{len}(s)}$
- $\textit{raw desirability}^3(P;\boldsymbol{c}) = \frac{\max_{v \in P} y(v;\boldsymbol{c})}{\sum_{s \in P} \text{len}(s)}$

The first sums the *average* values of venue interestingness on each path, the second sums the maximum venue interestingness on each *segment*, and the third sums the maximum venue interestingness on the entire *path*. Despite the significant differences among those definitions, we will find that the results do not change in any statistically significant way.

Finally, by considering the original definition of raw desirability, critics might rightly say that the evaluation results would be difficult to interpret. Ideally, the results are best interpreted in a comparative fashion. Hence, to ease illustration, we introduce the ideal recommendation strategy, call it *oracle*, and compare our method's desirability with the oracle's. The oracle is granted with the ability to foresee the actual (not predicted) interestingness of all the venues for every context. In each context $\boldsymbol{c}$, the venues are ordered according to their observed interestingness $y(v;\boldsymbol{c})$. Having this ideal recommender at hand, we can rephrase the definition of raw desirability with:

$$desirability(P;\boldsymbol{c}) = \frac{\textit{raw desirability}_{method}(P;\boldsymbol{c})}{\textit{raw desirability}_{oracle}(P;\boldsymbol{c})}$$

This quantity is easy to interpret as it will always be below 1 (no method can be better than the ideal one), and the closer to 1, the better the method's performance.

To further expand our comparative study, we introduce yet another class of recommendations: the *most interesting* path. This strategy takes the path with the highest cumulative interestingness $y(v) = \sum_{c \in C} y(v;\boldsymbol{c})$). The cumulative interestingness is computed on all the path's venues across all possible contexts. As a result, those recommendations are context-agnostic, in that, between two points, the same path is recommended no matter whether it is Friday night or Monday morning.

4.1.3 Validation Execution

Our experiments unfolds by recording desirability scores for a sample of 300 pairs of starting and ending locations in the whole area of Barcelona. Focusing on the use case of pedestrian navigation, we set the maximal geographical distance to be 1500 meters between each pair of locations[4], roughly less than 40 minute's walk. For each pair of destinations, a set of paths are generated. These consist of the shortest path, the most interesting path, and our own contextual path. We record the contextual $desirability(P;\boldsymbol{c})$ for those paths, and do so with a carefully selected validation methodology. To see why, consider that a traditional cross-validation would result in an experimental setup biased towards surprising good results. That is because, in the training phase, the algorithm that predicts the venue interestingness scores sees only the venues that have positive scores and does not see the venues that have zero scores (even though the latter are the majority). As such, no matter how wrongly the scores are predicted, as a by-product of cross-validation, the path recommender would still end up with routes that go through the observed venues and avoid the unseen ones, creating an artificial bias in the evaluation. By contrast, to have a validation that actually tests the effectiveness of our venue interestingness predictions, we should take a conservative approach. We do so by running our tests in a cold-start situation (cold-start as for context). That is:

- We select a subset of contexts (i.e., test set C^{test}) from the set of observed context C. For example:

$$C^{test} = \{\boldsymbol{c} = (h,w,d) | h \in \{11, 14-15, 19-23\},$$
$$w \in \{\text{Clear, Rain , Partly Cloudy}\},\ d \in \{\text{Tue, Fri, Sat}\}\}$$

- At each iteration, we leave one context $\boldsymbol{c} \in C^{test}$ out, and train the FM model for the remaining contexts $C \backslash \boldsymbol{c}$. We find that 30 is the number of latent factors for which AUC is maximum (Table 2). We then sample 300 pairs of destinations $\{(s,t)\}_{300}$; generate the shortest path, our contextual path, and the oracle path; record and compare those paths' $desirability(P_{s,t};\boldsymbol{c})$ scores.

In our experiments, we ignore the user-personalization component, by setting a $\alpha_u = 1$ and thus considering only the contextual cost; we set the strolling coefficient λ_c by line search over the interestingness distribution such that the extra-walking distance does not exceed 500 meters.

[2]www.wunderground.com

[3]Rain = {Light Freezing Rain, Thunderstorms and Rain, Heavy Rain, Rain Showers, storms and Rain}; Snow= {Heavy Snow, Snow,}; Fog ={Shallow Fog, Fog, Patches of Fog}; Overcast= { Drizzle, Light Drizzle, Light rain, }; Rain; Mostly cloudy; Clear; Scattered Clouds; Partly Cloudy and Unknown.

[4]This is the length of the straight line segment connecting two locations. The actual walking distance in the street might be slightly longer.

Table 2: Our prediction model's AUC for different numbers of latent factors K.

K	AUC	learning time(secs)
5	0.67	265.54
10	0.68	267.04
20	0.69	268.54
30	**0.692**	269.80
50	0.667	269.28
70	0.676	277.16
100	0.677	280.84

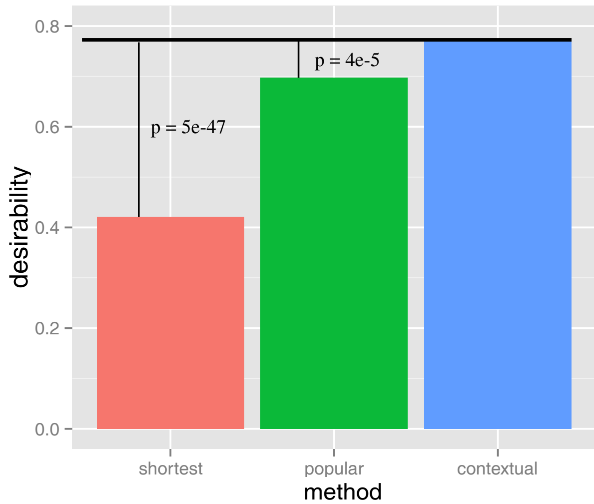

Figure 6: The mean desirability scores for shortest, most interesting (popular), and context-sensitive paths (compared to oracle paths). We test whether the three means differ statistically and find that they indeed do (the corresponding t-test's p-values are shown).

4.1.4 Recommendation Desirability

Figure 6 shows the mean desirability for the different paths. On average, our contextual paths show a 80% increase over the shortest paths (significance level $< 10^{-46}$) and 8.5% increase over the most interesting path (significance level of $< 10^{-4}$).

The difference of 8.5% is a promising - it reflects the ability of the model to return predictions for venue-context pairs unseen at training phase. Yet it is conservative result as it might reflect more the nature of the data than algorithmic performance. To see why, consider that only few venues happen to be sensitive to contextual changes, while many are not. Our results include all venues and, as such, are conservatively biased: as a thought experiment, imagine an extreme case in which a tiny fraction of venues were to be context-sensitive; we would not see any difference between a context-aware recommender and a context-agnostic one. However, that result would be an artifact of the data and would not speak to the actual context-aware recommender's performance. To test whether the difference between the contextual path's desirability and the interesting path's is limited by the abundance of contextual-insensitive venues, we try to gradually filter those venues out and see whether the mean desirability scores would consequently change. That is, we:

- Stratify venues based on their contextual sensitivity:

$$\text{sensitivity}(v;\boldsymbol{c}) = \left| y(v;\boldsymbol{c}) - \frac{1}{|C|-1}\sum_{\boldsymbol{c}_i \in C-\boldsymbol{c}} y(v;\boldsymbol{c}_i) \right|.$$

- Compute the desirability scores considering *only* the top $k\%$ most sensitive venues:

$$desirability_k(P;\boldsymbol{c}) = \frac{\sum_{y(v;\boldsymbol{c})>0 \wedge \text{sensitivity}(v;\boldsymbol{c})>q_k} y(v;\boldsymbol{c})}{\sum_{s\in P}\text{len}(s)},$$

where q_k denotes the k^{th} upper quantile of the sample {sensitivity(v;**c**)|y(v;**c**)>0} ($k = 100$ corresponds to considering all venues and results into the original $desirability$ values).

As one expects, we find that the more context-sensitive the venues at hand, the larger the observed desirability gap between contextual paths and interesting ones (Figure 7). That is simply because there is a large room for improvement, and contextual paths is able to capitalize on it.

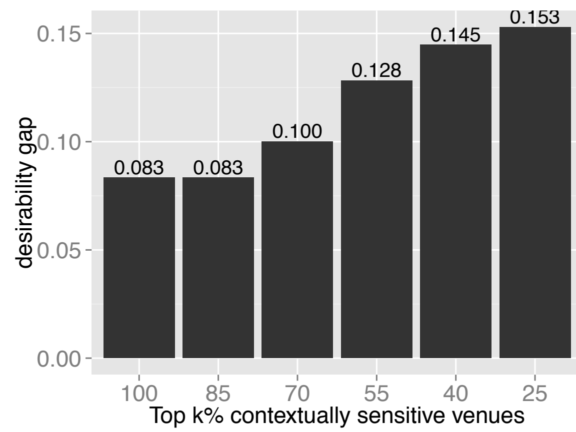

Figure 7: Difference between the contextual path's mean desirability and the most interesting (popular) path's. This difference increases as only the top k% most contextually-sensitive venues are considered.

4.1.5 Walking Overhead

So far we have ascertained that our framework does effectively suggest contextually relevant paths. Ideally, those benefits should not come at the price of lengthy routes. To see how much longer our recommendations take to go much more contextual-relevant paths, we compute their extra lengths (over shortest paths):

$$\textit{extra length} = \frac{\text{length(contextual)-length(shortest)}}{\text{length(shortest)}}.$$

By re-running our experiments for destinations at different distances and binning the corresponding results, we find that the contextual paths do not take much longer (Figure 8). Extra walking length is very limited for all destinations: for destinations at 500m, the contextual path requires 60 meter (1 minute) longer routes; for those at 2km, it requires 104 meter (2 minute) longer routes. On average, the median extra length is less than 6%.

4.1.6 Area Density

One might now wonder to which extent positive results are attainable in low-density areas of the city as well (e.g., outside the town historical center). To test whether spatial factors might influence our results, we divide the city into cells and compute the mean desirability score in each cell. More specifically, we pin down a grid of points on the map, where the points are equally spaced by 1500 meters; centered around those points, we draw squared cells with edge length of 4400 meters. This procedure results into a set of overlapping cells that uniformly cover the whole city. We then consider three distinct spatial factors for each cell as proxies for density: the number of street segments, of unique venues,

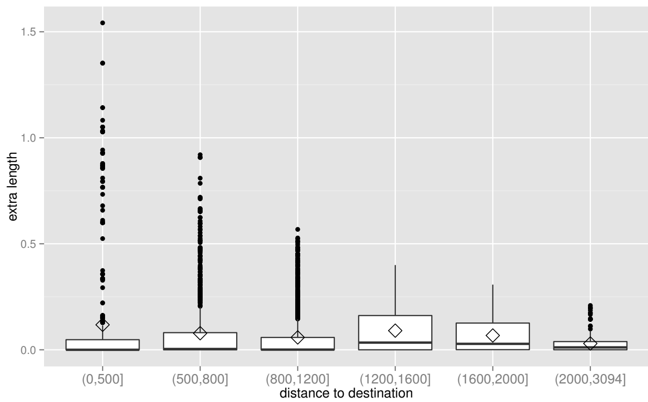

Figure 8: Our contextual path's extra walking length (over the shortest one) for destinations at increasing distance.

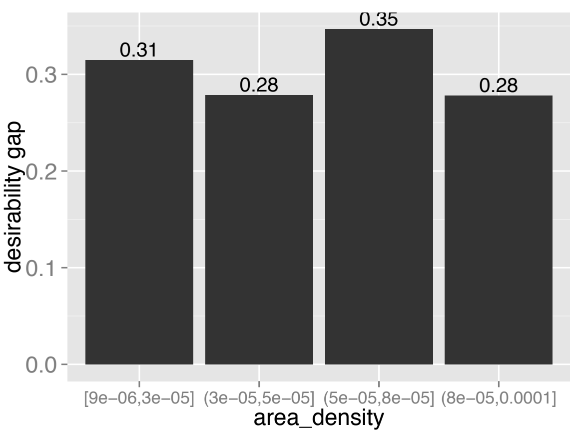

Figure 9: Difference between the contextual path's mean desirability and the shortest path's for areas of different density. By comparing the results corresponding to the different bars, we find that they are not statistically different: area density does not affect the desirability gap.

and of check-ins per square meters. Since those three measures are strongly correlated (with Pearson coefficients well over 0.85), we report the results for venue density. Remarkably, we find that mean desirability does not change with area density (Figure 9), suggesting that there is room for finding contextually-relevant paths in low-density neighborhoods as well.

4.1.7 Overheads

Our framework's storage overhead is minimal: loading the entire graph of the city takes less than 8MB and 4.1 seconds in running time[5]. The most demanding operations include training and generating the venue interestingness predictions (265 seconds), loading the transition matrices (4.1 seconds), and computing the contextual cost on each segment (12.9 seconds). However, those operations can be performed offline, and can be done only once. By contrast, each contextual path should be generated on-the-fly for each user request: that is indeed possible as generating a path takes as low as 5.8 milliseconds, on average.

[5]The results are obtained on a desktop machine with these specifications - CPU: Intel i5-3210 4-Core@2.50GHz, RAM: 2GB

5. DISCUSSION

User Study. Our evaluation has been quantitative and has not focused on whether users would find our proposal effective. To that end, as part of future work, we have conducted a preliminary user study involving 15 participants in Barcelona. Each participant is shown four paths between Arc de Triump and Barceloneta. The participants are asked whether they are familiar with the paths (i.e., whether they have visited the paths' area more than once and clearly remember the area's offering). The four paths are: a contextual path at 2pm (day path), a contextual path at 11pm (night path), the most popular path, and the shortest path. We choose 2pm because that is when lunch in Barcelona ends and coffee break follows, and 11pm because, at that time, dinner ends and bars and nightlife typically follow. The participants are asked to evaluate two aspects for each path. The first is how interesting each path is (on a Likert scale) for activities to be done at 2pm and for activities to be done at 11pm (we do not tell them which path is which). More specifically, the participants tell us whether they would take the suggested path at 2pm (or 11pm), and express their preference on a scale that goes from 1 (I would not take it at all) to 5 (I would definitely take it). The second evaluation aspect has to do with the extent to which the proposed path's length is acceptable: the score goes, again, from 1 (totally unacceptable) to 5 (it's the shortest path). Among our participants, the percentage of male-female is 62%-38%. The most common age band is that of 25-30. As for familiarity with the area, all our respondents have lived in Barcelona for more than a year and have visited the area under study more than once. All the participants find our proposal's paths to be of acceptable lengths (median 4, where 5 means no perceived difference between our paths and the shortest ones). They also score the day paths and the night paths highest in the corresponding day and night contexts (with median 4 on the scale [1,5] for both cases).

Tables 3 and 4 summarize the results by reporting the probability of observing a certain score for each type of path. For each of the two situations (2pm, 11pm), we report the participants' evaluations about the interestingness and acceptability of a variety of paths: the shortest path; "ours(2pm)" and "our(11pm)", which are the paths generated by our system at 2pm and 11pm; "popular", which is the path that maximizes the popularity of the venues across it (context-agnostic); and "long(2pm)" and "long(11pm)", which are the contextual paths generated without penalizing for length ($\gamma = 0$). Based on the overall scores (reported on the two rows $E[Interestingness]$ and $E[Acceptability]$), we see that the contextual paths are far more interesting than the other paths, and that they require a minimal extra walking distance (i.e., its walking overhead is always second to the shortest path). At both 2pm and 11pm, the contextual paths have higher expected interestingness than the other paths.

Table 3: Summary of results for paths at 2pm

	score	*shortest*	*ours(2pm)*	*ours(11pm)*	*popular*	*long(2pm)*
Interestingness	1	0	0	0	0	0.33
	2	0.13	0.2	0.13	0.2	0.07
	3	0.4	0.2	0.27	0.2	0.27
	4	0.33	0.27	0.4	0.53	0.07
	5	0.13	0.33	0.2	0.07	0.27
$E[Interestingness]$		3.43	**3.73**	3.67	3.47	2.91
Acceptability	1	0	0	0	0.07	0.93
	2	0.07	0	0	0	0
	3	0	0.07	0.27	0	0
	4	0.2	0.4	0.33	0.47	0
	5	0.73	0.53	0.4	0.47	0.07
$E[Acceptability]$		**4.59**	4.46	4.13	4.3	1.28

Interestingness from check-ins data. We have not shown any result for notions of interestingness other than venue popularity, not

Table 4: Summary of results for paths at 11pm

	score	*shortest*	*ours(2pm)*	*ours(11pm)*	*popular*	*long(11pm)*
Interestingness	1	0	0	0	0	0.33
	2	0.13	0.27	0.07	0.13	0.13
	3	0.4	0.27	0.47	0.4	0.27
	4	0.33	0.2	0.27	0.4	0.07
	5	0.13	0.27	0.2	0.07	0.2
$E[Interestingness]$		3.43	3.5	**3.63**	3.41	2.68
Acceptability	1	0	0.07	0	0	0.87
	2	0	0	0	0.07	0
	3	0	0	0.07	0	0.13
	4	0.27	0.4	0.27	0.4	0
	5	0.73	0.53	0.67	0.53	0
$E[Acceptability]$		**4.73**	4.32	4.64	4.39	1.26

least because of lack of data. Since we have proposed a framework (i.e., a module with plug-and-play components), one could plug a different "desirability component" in the future (this simply translates into having different values for $y(v_i; \mathbf{c})$).

Classes of Applications. Most of the past work has focused on navigation systems for tourists. Our model is able to generate different paths for tourists (by setting the familiarity coefficient α_u to a low value) and residents (by setting it to a high value) alike. This coefficient captures the idea that, while walking in a neighborhood, those not familiar with it rely on visible landmarks and avoid frequent turns (i.e., they do minimize angular changes), while residents are likely to know short-cuts and are comfortable with frequent turns that end up minimizing metric distance.

Modeling of Spatial Layout. Previous navigation approaches have modeled space in a coarse-grained way. Typically, the city layout is divided into cells, and those cells will then be nodes in the city's navigation graph. With such a representation, back-tracking and U-turns are not uncommon issues. To partly fix those issues, we have departed from this traditional spatial representation and have used space syntax. It is a simple graphical method that describes the way the different parts of space (e.g., street segments, entire areas) are connected to one another, mimicking the way individuals mentally deal with space and navigation. We are stressing the importance of researching alternative spatial representations not to unfairly criticize past work (after all, its focus was not on spatial modeling) but to stimulate further research in this very direction.

Engineering aspects. In the future, it might be also worth exploring two main engineering aspects. The first is the way we have implemented the axial map. Our current implementation uses OpenStreetMap's line representation, which is able to deal with curves but has some imperfections. As such, alternative approaches of implementing the axial map (e.g., raster analysis) might be explored. The second aspect has to do with the predictions of venue popularity. The factorization machine currently neglects the interactions between the variables (collinearity). Future work that will focus on venue popularity predictions might initially try a polynomial regression, or build a regressor that deals with high order.

6. CONCLUSION

The ethnographic observations made by urban sociologists such as Jane Jacobs, William Whyte, and Jan Gehl have resulted in workable principles about how people view, understand, and use city spaces [6, 13, 36]. These principles have drawn the attention of those who design cities but not of those who design navigation tools. Our framework has showed how to systematically apply some of those principles to the problem of automatic wayfinding, making important connections to the psychology of spatial cognition. In addition to principled spatial representations, one also has to consider contextual changes. Within the relatively unchanging spatial containers of a city, activities shift cyclically, and our framework accounts for contextual changes with state-of-the-art prediction techniques. To test the extent to which our framework is effective in suggesting personalized recommendations, we evaluated our proposal on real-data (by relying on Foursquare check-ins as a measure of popularity) and by conducting a preliminary user study in the city of Barcelona. Both evaluation suggest that the system is effectively able to recommend contextually-relevant paths that are only slightly longer than the shortest ones.

Our proposal could be readily extended in two main ways. First, our measure of contextual cost should be further validated to ascertain whether it actually reflects people's preferences. Second, other proxies for venue interestingness should be explored.

7. REFERENCES

[1] G. Adomavicius and A. Tuzhilin. Context-aware recommender systems. In *Recommender systems handbook*, pages 217–253. Springer, 2011.

[2] Z. Cheng, J. Caverlee, K. Lee, and D. Z. Sui. Exploring millions of footprints in location sharing services. *ICWSM*, 2011:81–88, 2011.

[3] C. Ellard. *You are here: Why we can find our way to the moon, but get lost in the mall*. Random House, 2009.

[4] R. J. Elliott and M. Lesk. Route Finding in Street Maps by Computers and People. In *AAAI*, pages 258–261, 1982.

[5] Z. Gantner, L. Drumond, C. Freudenthaler, S. Rendle, and L. Schmidt-Thieme. Learning attribute-to-feature mappings for cold-start recommendations. In *ICDM*, pages 176–185, 2010.

[6] J. Gehl and R. Rogers. *Cities for People*. Island Press, 2013.

[7] A. Gionis, T. Lappas, K. Pelechrinis, and E. Terzi. Customized tour recommendations in urban areas. In *WSDM*, pages 313–322, 2014.

[8] P. E. Hart, N. J. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. *Systems Science and Cybernetics, IEEE Transactions on*, 4(2):100–107, 1968.

[9] B. Hillier and J. Hanson. *The social logic of space*. Cambridge University Press, 1984.

[10] B. Hillier, A. Penn, J. Hanson, T. Grajewski, and J. Xu. Natural movement-or, configuration and attraction in urban pedestrian movement. *Environment and Planning B*, 20(1):29–66, 1993.

[11] H. Hochmair and A. U. Frank. Influence of estimation errors on wayfinding-decisions in unknown street networks - analyzing the least-angle strategy. *Spatial Cognition and Computation*, 2(4):283–313, 2000.

[12] C. Hölscher and M. Brösamle. Capturing indoor wayfinding strategies and differences in spatial knowledge with space syntax. In *Proceedings of 6th International Space Syntax Symposium*, pages 12–15, 2007.

[13] J. Jacobs. *The Death and Life of Great American Cities*. Vintage Books, 1961.

[14] A. Karatzoglou, X. Amatriain, L. Baltrunas, and N. Oliver. Multiverse recommendation: n-dimensional tensor factorization for context-aware collaborative filtering. In *RecSys*, pages 79–86, 2010.

[15] C. Li and K. Willis. Modeling Context-Aware Interaction for Wayfinding Using Mobile Devices. In *MobileHCI*, 2006.

[16] E. H.-C. Lu, C.-Y. Chen, and V. S. Tseng. Personalized trip recommendation with multiple constraints by mining user check-in behaviors. In *SIGSPATIAL*, pages 209–218, 2012.

[17] K. Lynch. *What Time is this Place?* MIT Press, 1976.

[18] D. R. Montello. Spatial orientation and the angularity of urban routes a field study. *Environment and Behavior*, 23(1):47–69, 1991.

[19] D. R. Montello and C. Sas. Human factors of wayfinding in navigation. *CRC Press/Taylor & Francis*, 2006.

[20] S. Möser, P. Degener, R. Wahl, and R. Klein. Context-Aware Terrain Visualization for Wayfinding and Navigation. *Computer Graphics Forum*, 2008.

[21] A. Penn and N. Dalton. The architecture of society: Stochastic simulation of urban movement. *UCL Press*, 1994.

[22] S. Porta, P. Crucitti, and V. Latora. The network analysis of urban streets: a dual approach. *Physica A: Statistical Mechanics and its Applications*, 369(2):853–866, 2006.

[23] D. Quercia, L. M. Aiello, R. Schifanella, and A. Davies. The digital life of walkable streets. In *Proceedings of the 24th International Conference on World Wide Web (WWW)*, 2015.

[24] D. Quercia, N. Ohare, and H. Cramer. Aesthetic capital: What makes london look beautiful, quiet, and happy? *CSCW*, 2013.

[25] D. Quercia, R. Schifanella, and L. M. Aiello. The shortest path to happiness: Recommending beautiful, quiet, and happy routes in the city. *Hypertext*, 2014.

[26] C. Ratti. Urban texture and space syntax: some inconsistencies. *Environment and Planning B: Planning and Design*, 31(4):487–499, 2004.

[27] S. Rendle. Factorization machines with libFM. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 3(3):57, 2012.

[28] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme. BPR: Bayesian personalized ranking from implicit feedback. In *AUAI*, pages 452–461, 2009.

[29] S. Rendle, Z. Gantner, C. Freudenthaler, and L. Schmidt-Thieme. Fast context-aware recommendations with factorization machines. In *SIGIR*, pages 635–644, 2011.

[30] S. Rendle and L. Schmidt-Thieme. Pairwise interaction tensor factorization for personalized tag recommendation. In *WSDM*, pages 81–90, 2010.

[31] K.-F. Richter. *Context-Specific Route Directions: Generation of Cognitively Motivated Wayfinding Instructions*. IOS Press, 2008.

[32] K.-F. Richter and M. Duckham. Simplest instructions: Finding easy-to-describe routes for navigation. In *GIScience*, pages 274–289.

[33] S. Scellato, A. Noulas, and C. Mascolo. Exploiting Place Features in Link Prediction on Location-based Social Networks. In *KDD*, pages 1046–1054, 2011.

[34] J. Shao, L. Kulik, and E. Tanin. Easiest-to-reach neighbor search. In *GIS*, pages 360–369, 2010.

[35] A. Turner. From axial to road-centre lines: a new representation for space syntax and a new model of route choice for transport network analysis. *Environment and Planning B: Planning and Design*, 34(3):539–555, 2007.

[36] W. Whyte. *The Social Life of Small Urban Spaces*. Project for Public Spaces Inc, 2001.