Computers, Environment and Urban Systems 73 (2019) 143–156

Contents lists available at ScienceDirect

Computers, Environment and Urban Systems

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

Where to go and what to do: Extracting leisure activity potentials from Web data on urban space

Demi van Weerdenburg[a], Simon Scheider[a,*], Benjamin Adams[b], Bas Spierings[a], Egbert van der Zee[a]

[a] Department of Human Geography and Spatial Planning, Utrecht University, the Netherlands
[b] Department of Geography, University of Canterbury, New Zealand

ARTICLE INFO

Keywords:
Place affordance
Urban space
Knowledge extraction
City planning
Latent semantics
Multi-label classification

ABSTRACT

Web data is the most prominent source of information for deciding where to go and what to do. Exploiting this source for geographic analysis, however, does not come without difficulties. First, in recent years, the amount and diversity of available Web information about urban space have exploded, and it is therefore increasingly difficult to overview and exploit. Second, the bulk of information is in an unstructured form which is difficult to process and interpret by computers. Third, semi-structured sources, such as Web rankings, geolocated tags, check-ins, or mobile sensor data, do not fully reflect the more subtle qualities of a place, including the particular functions that make it attractive. In this article, we explore a method to capture leisure activity potentials from Web data on urban space using semantic topic models. We test three supervised multi-label machine learning strategies exploiting geolocated webtexts and place tags to estimate whether a given type of leisure activity is afforded or not. We train and validate these models on a manually curated dataset labeled with leisure ontology classes for the city of Zwolle, and discuss their potential for urban leisure and tourism research and related city policies and planning. We found that multi-label affordance estimation is not straightforward but can be made to work using both official webtexts and user-generated content on a medium semantic level. This opens up new opportunities for data-driven approaches to urban leisure and tourism studies.

1. Introduction

As part of the smart city wave and corresponding data-driven developments (Townsend, 2013), the value of Web data for computable place representations is finding its way into urban leisure and tourism studies (Marine-Roig & Clavé, 2015) and related city policies and planning (Kitchin, 2014). Since leisure plays an increasingly important role in today's urban economy (Lorentzen, 2009), many cities have adopted data-driven strategies for planning and promoting an attractive city. This includes, e.g., the city of Zwolle, a Hanseatic city of about 125.000 inhabitants in the northeastern part of the Netherlands. Making the old town of Zwolle the "place to be" in the Northeast is the ambition of the urban authorities for the near future. To achieve this, local planners are exploring data driven strategies to improve the information provision towards potential visitors. An example of the present plans is to create a dynamic digital map showing what is happening where in the city centre of Zwolle (Bureau voor Economische Argumentatie, 2017). In order to recommend places for city marketing and to monitor and guide visitor flows in the city centre, data of high spatial, temporal, and semantic resolution is needed (Batty, 2013; Buhalis & Amaranggana, 2013). This includes not only the numbers of facilities and flows of visitors (so called "hard data"), but also perceived and experienced qualities of the many places that tourists and residents may visit to experience a city (so called "soft data").

While the former kind of data is nowadays easy to gather, the latter, though called "soft", is much more difficult to obtain. Looking at popular websites such as of the Zwolle Tourist Agency (VVV) and TripAdvisor, e.g., quickly and easily reveals the Museum de Fundatie (Fig. 1a) and the Sint-Michaëls church (Fig. 1b) as two of the main tourist highlights. However, the more subtle functions of these places often remain invisible. Who would have guessed that the museum is actually a nice place to go for a drink and that you can buy books in the Sint-Michaëls church (cf. Section. 6.2)?

Indeed, "soft data" is very hard to obtain in a manner that is both *scalable* and *reliable* (Batty, 2013). Scalable means that information can be obtained in large quantities and for many cities. Reliable means that this can be done in a way that makes the subtlety of place qualities discoverable for residents and visitors and exploitable for city

* Corresponding author.
E-mail address: s.scheider@uu.nl (S. Scheider).

https://doi.org/10.1016/j.compenvurbsys.2018.09.005
Received 15 May 2018; Received in revised form 22 August 2018; Accepted 23 September 2018
Available online 05 October 2018

(a) Museum De Fundatie, Zwolle. Source: Bierman Henket architecten, foto: Joep Jacobs. CC BY 2.0

(b) Sint-Michaëlskerk, Zwolle. Cornelis Springer, public domain.

Fig. 1. Two famous touristic places in Zwolle.

researchers, consultants, policy makers, and planners. While the Web in principle offers such information on places (Ballatore, Wilson, & Bertolotto, 2013) stemming from websites of municipalities, service providers, and geolocated social media posts (Cataldi, Ballatore, Tiddi, & Aufaure, 2013; Goodchild, 2007), the difficulty lies in extracting it in high fidelity for Geographic Information Systems (GIS) (Jonietz, 2016; Scheider & Janowicz, 2014). Even though the number of studies applying Web-data analysis in the field of tourism and leisure is increasing, the majority of these studies still apply manual processing and analysis of the data (Lu & Stepchenkova, 2015). To overcome the problems of scalability and reliability, useful methods have been suggested in the field of natural language processing (NLP) and data mining (see e.g. Alazzawi, Abdelmoty, & Jones (2012)). One option is to use topic models such as Latent Dirichlet Allocation (LDA) (Blei, Ng, & Jordan, 2003), which postulates topics as hidden graphical nodes to generate probabilities over words (Blei, 2012), making it an important discovery method for the Semantic Web (Gangemi, 2013).

So far, topic models have primarily been used to detect urban patterns in an unsupervised manner and for categorizing places (compare Section 2). Extracting information on what activities a place actually affords (*place affordance modeling* (Jordan, Raubal, Gartrell, & Egenhofer, 1998; Scheider & Janowicz, 2014)), however, is a *supervised multi-label classification* problem. Activity classes must be made explicit in order to be shared and compared with other empirical studies (e.g., traditional surveys and activity diaries (Crosbie, 2006)), and therefore the learning process must be supervised. Furthermore, even though places are usually tagged with a single category ("café"), they often afford multiple kinds of activities at the same time, some of which might be unexpected or not explicit in data, and therefore currently lacking in Web based geographic research (Ballatore, 2014). Solving this problem would enable high resolution maps of urban leisure potentials, imputing a qualitative dimension to quantitative measurements.

To approach these challenges we investigate the following questions in the local context of the city of Zwolle:

1. How and how reliably can we extract place-bound leisure activity potentials from different Web sources in an automated fashion?
2. How can we study urban leisure potentials across a city like Zwolle, taking into account side activities at infrequently visited places?
3. What role could these methods play in urban leisure and tourism research and related city policies and planning?

In Section 2, we situate our work into both urban studies and information science. In Section 3, we introduce our methodology and an ontology of leisure activities. In Section 4, we present our methods for Web data preprocessing, and then test three supervised multi-label learning strategies on a manually labeled dataset in Section 5. Results are discussed in Section 6 with respect to all three questions.

2. Urban leisure, web data and semantic modeling

In this section, we embed our article in the existing work on urban leisure and tourism, as well as on Web data and semantic models.

2.1. Urban leisure and the visitor economy

Cities, their historical cores in particular, have long been appealing to people with a craving for activities and experiences related with culture, leisure, and consumption (Gospodini, 2006). However, it is only since the end of the 20th century that this appeal is being planned and promoted with the deliberate aim to build a 'visitor economy' (Barrett & Hall, 2018). Such a visitor economy relies on a multi-functional mix of facilities together with visually attractive environments in the city (Zukin, 1998). In an attempt to attract a large number and variety of visitors, cities redevelop and market themselves as "must visit" destinations offering fun and enjoyment (Spierings, 2006). Themed experiences, combining culture, leisure and consumption, are (re)imagined and allocated to specific urban spaces to create a 'theme park' city (Hannigan, 1998), including e.g., 'historic city', 'cultural city', 'sport city', or 'nightlife city' (Ashworth & Page, 2011). It is within the historical cores of cities that such themed experiences are linked through a network of walking routes with the aim to extend visiting time, to increase exposure of the amenities on offer and to seduce visitors into spending money (Spierings, 2013). However, as opposed to top-down planned and managed theme parks like Disneyland, city centers are living entities in which an evolving amalgam of service providers and attractions together form dynamic consumption spaces (van der Zee, van der Borg, & Vanneste, 2017). This makes it much more challenging for policy makers and planners to manage and market the attractiveness of city centers, and for residents and visitors to overview and get accurate information on the leisure affordances on offer.

2.2. Web data about urban leisure

Information about urban leisure is being circulated and marketed through a diversity of media–including international branding campaigns, local newspapers, tourist maps and brochures, websites, apps, and social media (Hanna & Rowley, 2015; Xiang & Gretzel, 2010). City visitors increasingly utilize Web information sources and social media to make decisions about 'where to go and what to do' (Hudson & Thal, 2013; Xiang & Gretzel, 2010). As a consequence, the growing amount of Web information has the potential to reveal time-space patterns and related activities (Girardin, Calabrese, Dal Fiore, Ratti, & Blat, 2008; Kwan, 2002; van der Zee, Bertocchi, & Vanneste, 2018), as well as the perceptions and experiences of leisure seekers (Lu & Stepchenkova, 2015; Marine-Roig & Clavé, 2015; Ye, Zhang, & Law, 2009).

However, qualitative information about places is dependent on perception (Montello, Goodchild, Gottsegen, & Fohl, 2003) and thus seldom utilized due to its vagueness and the difficulty of extracting it in a reliable way. In fact, researchers mostly draw on structured "hard" information, such as geo-coordinates, tags, ratings, or rankings derived from a single user generated source (such as place review sites or photo-sharing platforms), and which are fairly easy to collect and process (Girardin et al., 2008; Hollenstein & Purves, 2010; Hu et al., 2015; Purves, Edwardes, & Wood, 2011; van der Zee et al., 2018). However, while visitors and inhabitants of a city leave them in abundance, these traces reflect qualitative information only to a limited degree. For one, geo-coordinates, tags and ratings require interpretation in context in order to reconstruct the actual experiences and functions of corresponding places (Cataldi et al., 2013; Jonietz, 2016; Scheider & Janowicz, 2014). Furthermore, social media data has been criticized for providing a biased representation of reality which is heavily influenced by a platform's algorithms (Caliskan, Bryson, & Narayanan, 2017; Scott & Orlikowski, 2012). For this reason, there is a danger of overlooking functions that are less known or less popular (false negatives).

2.3. Semantic modeling of urban leisure

To overcome these problems, a different approach is needed that makes use of "soft" information about places from a diversity of sources. The latter is mostly available in an *unstructured (textual) or semi-structured form* (Adams & Janowicz, 2015; Purves & Derungs, 2015) on websites as well as from social media, requiring human text interpretation (Lu & Stepchenkova, 2015). Text corpora can be analyzed according to semantic and geographic dimensions (Fabrikant & Buttenfield, 2001), yet it is an open problem how semantic dimensions of places can be usefully extracted. Some authors have utilized text mining approaches on content in order to account for the more qualitative experiential qualities of space, such as thematic place signatures (Adams & Janowicz, 2015; Adams & McKenzie, 2013; McKenzie, Janowicz, Gao, Yang, & Hu, 2015), sentiment analysis (Cataldi et al., 2013; Guo, Barnes, & Jia, 2017; Ye et al., 2009), and sound experiences (Chesnokova & Purves, 2018).

These approaches mostly rely on unsupervised learning of ad-hoc activity concepts, denoting, e.g., linguistic (Alazzawi et al., 2012) or behavioral (Adams & Janowicz, 2015; Farrahi & Gatica-Perez, 2011; Hasan & Ukkusuri, 2014) patterns, semantic signatures (McKenzie et al., 2015) or cluster membership (Hu et al., 2015). Unsupervised labels, however, lose their meaning outside of the context of the data set on which they were learned.[1] Furthermore, manual qualitative analysis of Web data is difficult to replicate and scale up (Lu & Stepchenkova, 2015). For this reason, research is needed that assesses the quality of supervised methods for urban leisure activity modeling. In such a model, labels obtain meaning in terms of a shareable ontology, which provides urban planners and policy makers with formal information on place affordances they are unable to retrieve manually.

[1] Teaching a machine a human concept, such as detecting an activity, requires "human supervision", i.e., manually curated and labeled training data (Friedman, Hastie, & Tibshirani, 2001).

3. Methodological preliminaries

In this section, we explain our methodology. In particular, we introduce an information ontology that can be used to automatically capture leisure affordances at urban places in an explicit rather than implicit way.

3.1. Approach

Progress towards a reusable and automated place affordance model requires, as a first step, creating an explicit, machine readable vocabulary about the various kinds of activities that can occur at urban places. This vocabulary serves to structure leisure space and links to more traditional leisure and tourism research on urban activities. For this purpose, we develop a leisure ontology in Section 3.2, compare Fig. 2. The ontology is then used to prepare training data (Section 4) used for training and testing various supervised classifier models (Section 5). These models are needed in order to *automate* the process of extracting leisure activity potentials from various web text sources and tags on urban places, scaling up affordance analysis across cities (Fig. 4, white boxes). Using cross-validation,[2] we measure the quality of our models in this respect and discuss to what extent future automated extraction is possible for other cities and different places (Section 6.1). We then discuss how urban leisure activity potentials can be analyzed across a city like Zwolle, using *affordance density maps* (Section 6.2), and what this means for future urban research (Section 6.3).

3.2. Urban leisure ontology

An information ontology is a "formal specification of a shared conceptualization" (Gangemi & Presutti, 2009). We formalized our ontology using the Web Ontology language (OWL),[3] a syntax standard of the World Wide Web that can be used to publish and share ontologies on the Web. We based our design on the *Place Activity* design pattern[4] (Scheider & Janowicz, 2014). In this pattern, place affordances are defined as *activities afforded* by *places*, and in which various *referents* can be involved (see Fig. 3), such as different types of food involved in eating at a restaurant. Places are (geo)*located* and activities *happen* at some time. Each of the circles in Fig. 3 denotes an OWL class, and each one of the arrows stands for some relation between instances of these classes. The arrow "are" says that place activity is a subclass of affordance.

This pattern was then filled with more specific leisure activity subclasses in an iterative manner (see Fig. 4) to obtain the *Urban Leisure ontology,*[5] abbreviated with the prefix *ulo*. To design relevant and reusable leisure classes, we made use of several scientific resources. We based our main leisure categories on the work of Ashworth and Page (Ashworth & Page, 2011), and then refined these with categories from published geographical studies including (Szalai, 1972), and the Dutch studies "Met het oog op the tijd" by SCP 2013 (Clon, 2013) and the factsheet "ContinueVrijeTijdsOnderzoek" by NBTC-NIPO (NBTC-NIPO Research, 2016). The latter studies consider leisure day trips and activities outside home. They do not include activities with an overnight stay (such as in a hotel) and visits of friends. We also used the websites of local places in Zwolle to cover locally important types of activities and search queries on the travel review website TripAdvisor.[6] During

[2] To simulate unknown samples (in this case places from unknown cities), our training sample is split iteratively into training and test data. For each subsample, the model is built on the training sample and tested on the test sample (Friedman et al., 2001).

[3] https://www.w3.org/OWL

[4] http://geographicknowledge.de/vocab/PlaceActivity.ttl

[5] *ulo*: `http://geographicknowledge.de/vocab/UrbanLeisure.ttl`

[6] `https://www.tripadvisor.nl/`

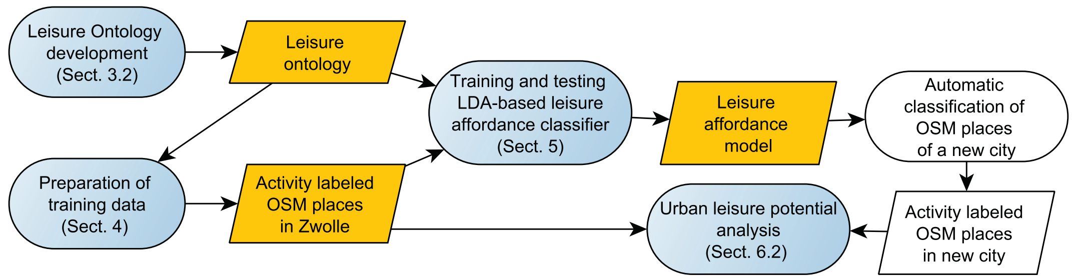

Fig. 2. Methodology followed in this article. Round boxes denote steps taken, parallelograms denote results. Empty boxes are outside the scope of this article.

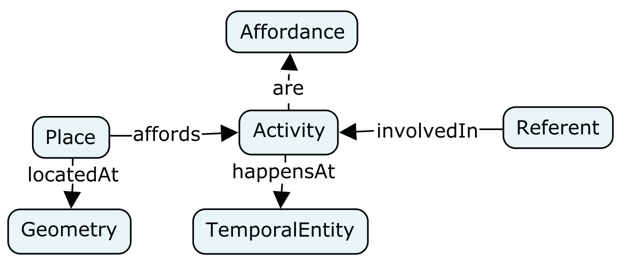

Fig. 3. Place activity design pattern.

the process of data preparation and modeling (see next section), we always came back to refine the ontology with missing classes. Our design pattern proved successful in the sense that it covers a large variety of activity types based on recombining referent and activity classes (see Figs. 4 and 5). In this way the same activity class, such as "Eating", can be specialized by different kinds of referents, such as "Lunch" and "Ice cream", and vice versa. Furthermore, it also allowed us to choose an appropriate *semantic level of detail* for modeling, by collapsing or expanding child nodes in the classification tree and testing learning models on that level. Though the ontology was developed for the urban

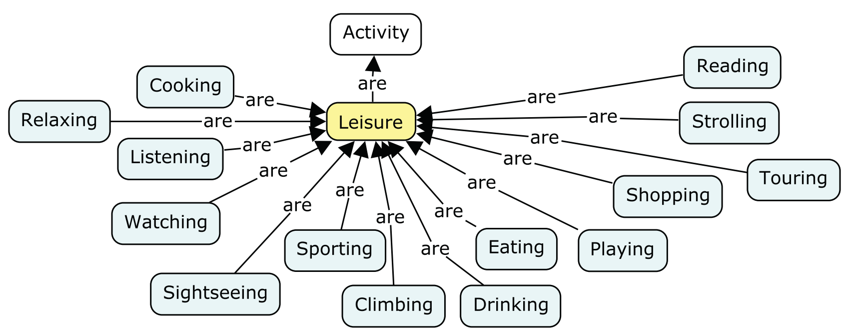

Fig. 4. Upper level of leisure activities in the ontology (level subsumed by ulo:Leisure).

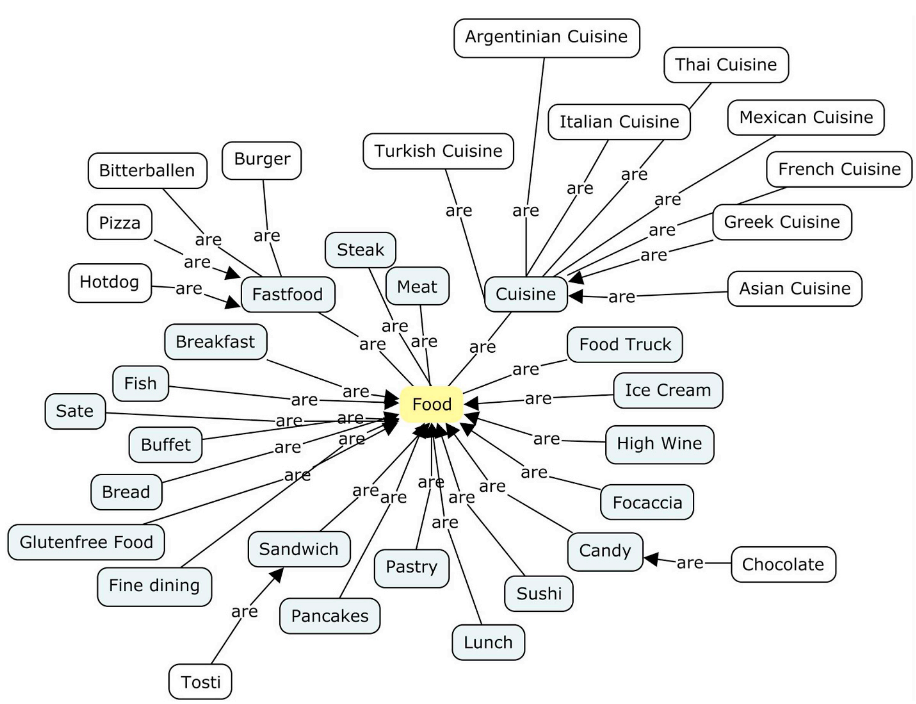

Fig. 5. Activity referent classes related to "Food".

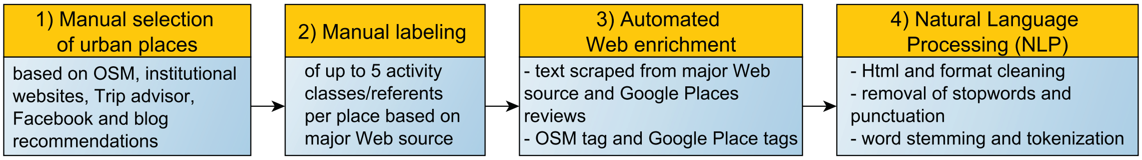

Fig. 6. Preprocessing steps for place sampling and Web enrichment.

settings in Zwolle, it can probably be used across similar middle-sized cities.

4. Preparing place data for learning from the Web

In this section, the Web data collection and preprocessing steps (compare Fig. 6) are described, which provided a training dataset suitable for affordance learning.

4.1. Manual selection and labeling of urban places

Places were selected with the following constraints:

1. Places should be uniquely described by a website having substantial activity-related content;
2. Places should be georeferenced and identified in Open Street Map (OSM).

The first restriction makes sure activities are place-bound. Many websites contain an overview of possible activities at multiple places (e.g., “the best addresses to shop in Zwolle”), and thus are not appropriate for our model. The second restriction enables us to cross-identify places, and thus to combine descriptions from different sources, as well as georeference and map them.

The encoding was performed during summer months from May through July 2017. Note that some websites were season-dependent, which explains that some activities such as sitting on a terrace or shopping fairs are frequently present in the training data, while, e.g., a winter fair is missing. Place information from a variety of sources was collected. First, places described in websites developed by local institutions were selected, for example, the official leisure website of the municipality and the official website of the Zwolle Tourist Agency (VVV). In addition, Google recommendations were used when searching for “Zwolle” or “Zwolle Nederland interesting places”, see Fig. 7. Second, user-generated sources were accessed, such as the reviewing website TripAdvisor, which at the moment are among the most popular information sources for tourists and leisure seekers (Lu & Stepchenkova, 2015; Marine-Roig & Clavé, 2015). Finally, as a third resource we retrieved sites in the blogosphere, for example, “In de Buurt Zwolle,” see Fig. 8. All posts in the months June and July on the subpages “to do” and “food & drinks” were retrieved.

Places described in these sources were identified as OSM nodes or, in the case of some streets, as OSM ways. OSM nodes and ways are labeled with tags, which were also used as part of the training data for the model. Fig. 9 shows how OSM mappers interpret places using tags such as ‘worship’, ‘shopping’, and ‘café’, visualized as symbols on the map.

Overall, information on 189 different venues[7] within the inner the city of Zwolle were collected and manually labeled based on their *currently* afforded activities. The frequency of activity labels is illustrated by Fig. 10. Many places offer multiple activities, such as a park or a bakery which also has a café and sells coffee. To limit the complexity of the model, we used a maximum of *five* kinds of activities per place, each including an activity and a referent class, resulting in 326 different affordances. Finally, the URI of the *main website* used for labeling was stored in the training data (Table 1).

4.2. Automated Web enrichment

To integrate further sources of structured and unstructured user-generated content into our data set, we automatically linked OSM objects with Google place objects based on both place name similarity and geographic distance, using Google's *text search engine*[8] with a maximal spatial radius of 300 m. Note that while the matching quality here is defined by the Google service, manual inspection showed a very high correspondence of matches using these two criteria.

Using the Google Places API,[9] we then retrieved *review texts* and all available *Google place tags*. Furthermore, via the Overpass API, we obtained corresponding *place tags* from *OSM.*[10]

The OSM tags (*keys*) deemed relevant in this context are listed in Table 2. For every such key, we retrieved the corresponding OSM *value* as a tag (e.g., “cuisine”: “japanese”), and encoded them as nominal variables. We then used Web scraping to obtain the text from the main website as described in the last section, using Richardson's “Beautiful Soup”[11] to clean up HTML and filter out script and style elements. The result of this enrichment is available online as a *json* file.[12]

4.3. Cleaning and tokenizing Web texts for NLP

To train an LDA model, an important first step is cleaning texts of punctuation and stop words which have minor semantic relevance. Furthermore, to identify words across grammatical varieties, such as inflections, it is necessary to *tokenize* words; i.e., to transform them into a “normal” form using the word stem. For these purposes, we used the stopword and punctuation lists of the NLTK[13] and Gensim[14] packages, including Porter (for English) and snowball (for Dutch) stemmers. The tokenized texts were then summarized into a *document-term matrix*, a matrix of word counts per text, which is input format to various supervised versions of LDA based learning, as described in the following section.

4.4. Reflections on the data quality

The data quality turned out to be heavily dependent on the web enrichment. From 189 places, only *153* were complete with respect to web scraped texts and tags. This was a result of varying *quality of service* while retrieving web texts. Furthermore, only 66 of these places were complete with respect to Google reviews, due to the many missing entries at Google.

The labeling captured 20 unique activity classes in Zwolle, including 105 different referent classes and 62 different sorts of places. As

[7] https://github.com/simonscheider/PlaceLDA/blob/master/Encoding010817.xlsx

[8] https://developers.google.com/places/web-service/search

[9] https://developers.google.com/places

[10] https://wiki.openstreetmap.org/wiki/Overpass_API

[11] https://www.crummy.com/software/BeautifulSoup/

[12] https://github.com/simonscheider/PlaceLDA/blob/master/training_train_u.json

[13] http://www.nltk.org/

[14] https://radimrehurek.com/gensim/

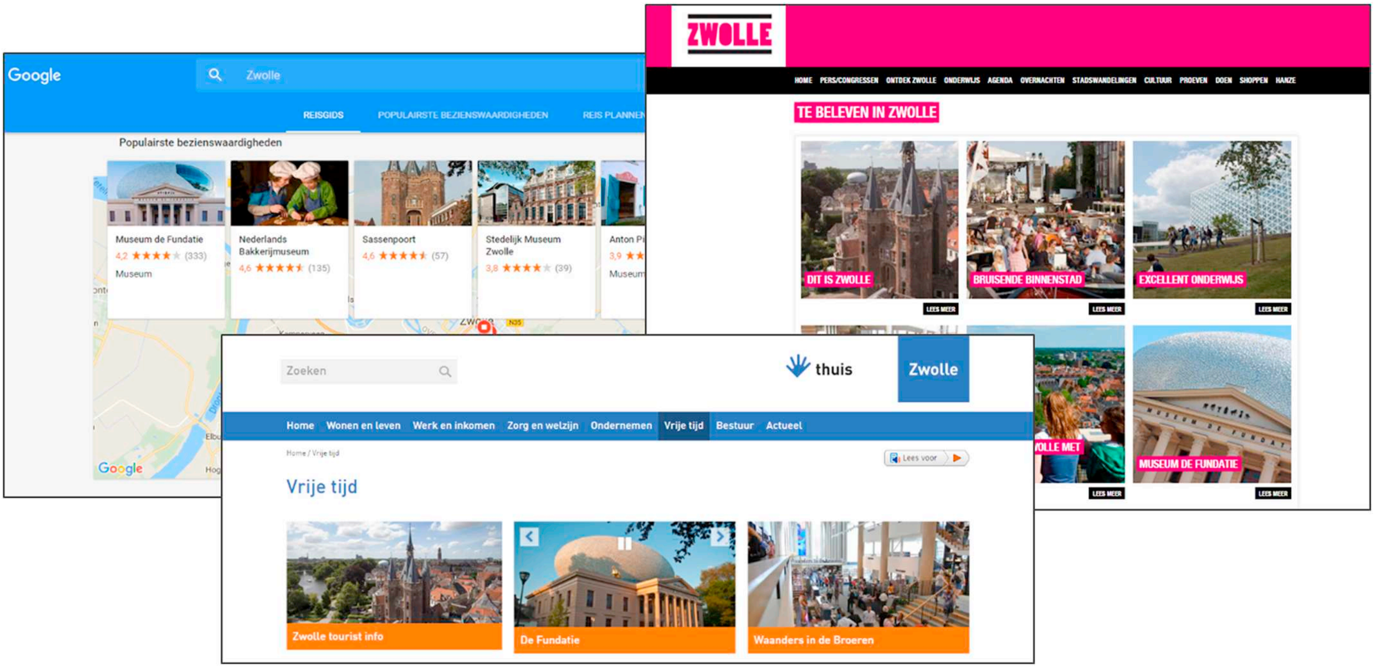

Fig. 7. Touristic webpages (http://zwolle.nl/vrije-tijd, http://zwolletouristinfo.nl) and Google recommendations for Zwolle.

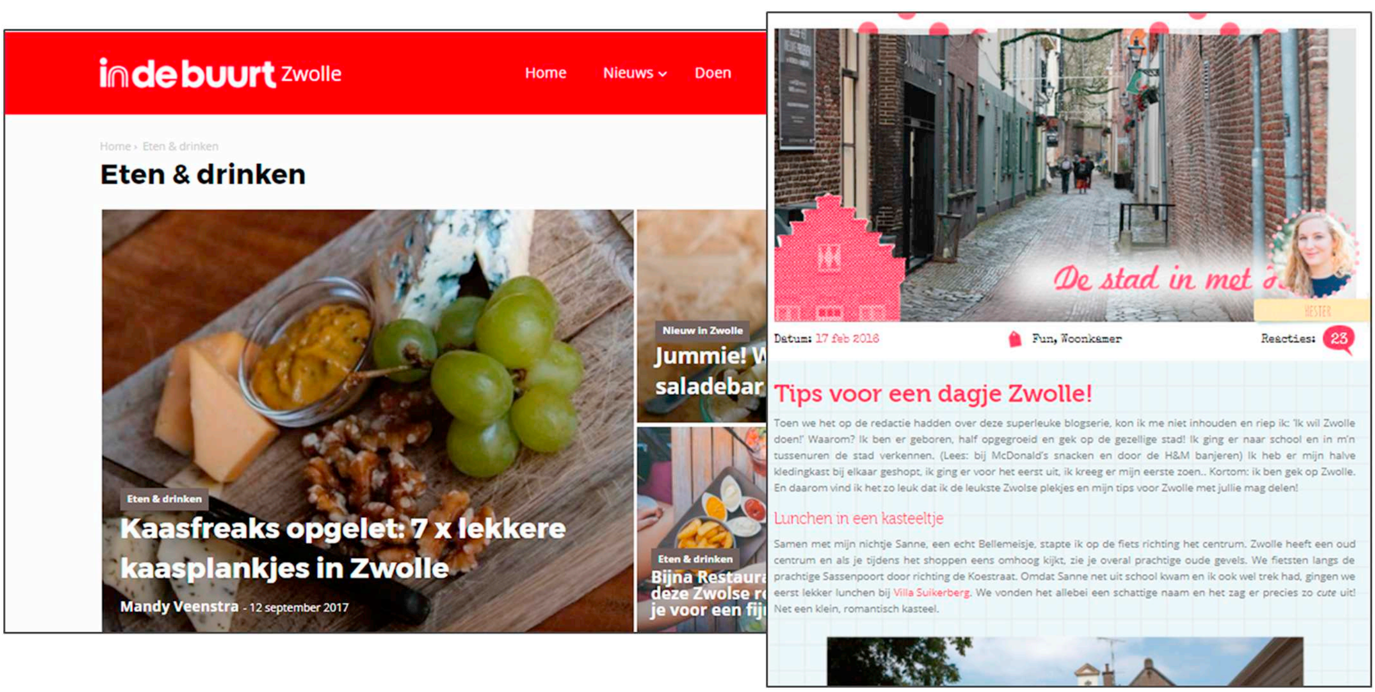

Fig. 8. Touristic blogs ("In de Buurt Zwolle", blog "Huis van Belle") for Zwolle.

can be seen in Fig. 12, the distribution of activities is skewed with a long tail, where two activities, namely Eating (29.8%) and Shopping (20.6%) account for half of all instances within the data set. This corresponds to the conclusions of the municipal report 'Buzzing Zwolle', which states that Zwolle is still largely focusing on shopping and food service.[15] The third and fourth important categories are Drinking (15.6%) and Watching (10.7%), where the latter is often related to referents such as architecture or an exhibition in a gallery.

A similar pattern occurs within the distribution of place classes, where restaurants and cafés, parks and squares dominate (see Table 3). The referent classes, however, are more dispersed. The ten most occurring referent classes make up only 41% of the total, and many of them are only encoded once or twice. As expected, food (7.4%) is the most frequent class, but architecture, terraces, gifts and exhibitions are among the 10 most frequent classes, too.

During data collection, it became clear that some activities that are typically conducted in urban areas and often described on web pages, such as 'shopping in the city' *are difficult to capture as a place in OSM* because they are not necessarily place-bound. Leisurely shopping

[15] https://www.zwolle.nl/sites/default/files/strategische-agenda-binnenstad-2017-samengevat.pdf

Fig. 9. Screenshot of the Grote Markt, Zwolle as shown on OSM.

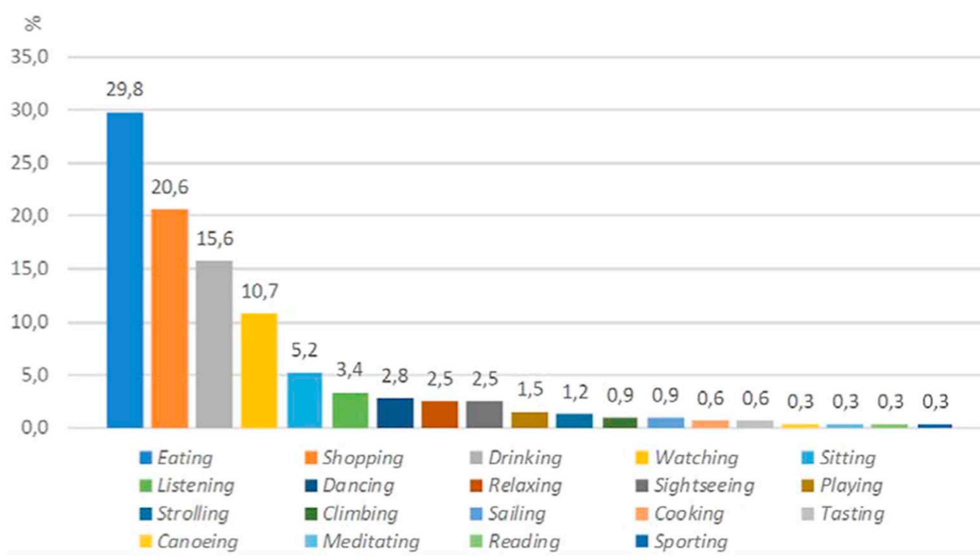

Fig. 10. Distribution of activity classes over the encoded places of the city center of Zwolle.

Table 1
Example of manual place encoding.

OSM identifier	Place name	ulo:Activity	ulo:Referent	ulo:Place	Website
- Node (osm)		As in ontology	As in ontology	As in ontology	URL
- Way (osmw)					
- Relation (osmr)					
osm:2500428169	Hedon	ulo:Listening	ulo:Concert	ulo:Theatre	https://.

Table 2
Keys of enriched training data file.

Key	Description
Class	*Activity class manually added in terms of ulo ontology, format ulo:Activity and ulo:Referent*
ulo_Place	*Place type manually added in terms of ulo ontology*
Website	*URL of the website used to scrape place descriptions*
Web title	*Title of the website used to scrape place descriptions*
webtext	*Text of the website used to scrape place descriptions (cleaned with Beautifulsoup)*
Name	*Name of the place (manually added)*
Review text	*Text of Google Places reviews (if available). Google place information was added based on the spatial distance and name similarity*
Google type	*Place tags from Google Places (if available). (in alphabetical order)*
Google ID	*Google Place ID (if available)*
Latitude	*WGS 84 Y Coordinate (taken from OSM, converted to centroid for ways)*
Longitude	*WGS 84 X Coordinate*
OSM keys	*Open Street Map keys with respective values, or 'No' if missing ['shop', 'amenity', 'leisure', 'tourism', 'historic', 'man_made', 'tower', 'cuisine', 'clothes', 'tower', 'beer', 'highway', 'surface', 'place', 'building']*

Table 3
Distribution of place and referent classes within the dataset.

#	Place	%_Place	Referent	%_Referent
1	Restaurant	23.9	Food	7.4
2	Cafe	13.2	Architecture	6.4
3	Park	4.0	Coffee	4.3
4	Square	4.0	Lunch	4.0
5	Bar	3.4	Terrace	4.0
6	Church	3.1	Gifts	3.4
7	Museum	3.1	Exhibition	3.1
8	Historic Building	2.8	Fashion	3.1
9	Fashion Store	2.5	Beer	2.8
10	Bookstore	1.8	Wine	2.8
Total		61.7%		41.1%

means strolling around the shopping streets without the goal of going to a specific shop (Spierings, 2006), similar to taking a walking tour or doing sightseeing from a boat or tourist vessel. Another issue that occured during the data collection is that not all sights in the city *have a website that describes an activity*. An example is the old city hall at the Church square, which is part of the municipality and lacks any parseable text.

5. Modeling leisure affordances with semantic topics and tags

In this section, three supervised machine learning approaches to model leisure affordances are described, which we refer to as *slLDA*, *mlLDA*, and *LLDA*. We test these approaches with different parameterizations, for which we report the estimation quality and discuss the results.

5.1. slLDA

For slLDA, the task was reduced to a *single-label* supervised learning problem in a pipeline with unsupervised LDA. We generated a class frequency distribution and selected the most frequent class for each place in the data set. Features for learning included both semantic topics extracted from the texts using unsupervised LDA (Blei, 2012) and the place tags identified during the enrichment process described in the previous section. This dataset was then used as training data with various ordinary (single-label and multi-class) machine learning classifiers. Using the most frequent class has the advantage that we can test the multitude of powerful ML classifiers that are available, yet has the drawback that activity variety about a place is lost. Furthermore, since LDA is unsupervised, topics irrelevant for classification may occur. This approach is illustrated in Fig. 11.

5.1.1. Model parametrization and feature selection

We varied the *single label multi-class classifier learning* algorithms using standard parameters, which were taken from the *scikit-learn* libary,[16] see Table 4 (the "Neural Net" algorithm was a *multilayer perceptron*). In addition, the text sources were varied using Dutch language web scraped texts and English Google reviews. We also varied our model with respect to whether it made use of OSM/Google place tags or not, and with respect to the *semantic depth* of the activity class label in our ontology. The latter was done by choosing either the general activity class, or synthesizing a combination of activity class and referent class.[17] Finally, we restricted models to those classes with at least five instances, in order to prevent overly sparse labels. For parametrization of the *LDA learning*[18] algorithm, we chose 18 topics learned with 600 iterations and a fixed random seed. Results remained stable when testing other numbers of topics.

5.1.2. Results and evaluation

The LDA model was trained with a *corpus vocabulary* of 7067 words occurring 34,326 times over 153 web scraped place description texts. It produces vectors of probabilities over words for each topic, which can be used to estimate the probability that a given topic occurs in a given text. Fig. 12 shows the 18 topics in terms of their 5 most probable words. Many of these topics can be well interpreted. For example, topic 5 with word stems "smak" (for "smaakt", "smaakvol", or "smakelijk", all related to 'taste') in combination with "bier", "caf", "wijn" can be interpreted as a place where tasteful beers and wines are served in a café. Topic 3 summarizes beer brewery concepts, while topic 0 is related to movies (A cinema in Zwolle is called "Pathé" with its own 'unlimited' cinema card), and Topic 16 is related to art galleries. With other topics (such as 8, 9, 10, and 14), interpretation becomes more difficult, since stemmed website artifacts (such as Web addresses) range among the words. Our conclusion is that in some cases, Web artifacts take over topics, due to our comparably small sample. The amount of such artifacts dramatically decreases with Google reviews. However, the review texts also seem to be more homogeneous and centered around food, which might be a problem when considering semantic coverage.

Regarding the *classification quality*, we measured the performance of each single-label classifier with *10-fold cross validation* and standard quality metrics (accuracy, weighted precision, weighted recall and F-measure (a combination of precision and recall)) and compared this against a *naive model*, which was based on the majority class. Regarding model parameters, it turned out that restricting the class size to ≥ 5 was most successful for all text sources and algorithms. Furthermore, trying to estimate activities down to the level of referent classes largely failed. In the following, results will therefore be reported for highest level activity classes and class size ≥ 5, if not stated otherwise.

As can be seen in Table 5, the classification problem is difficult, with some classifiers (Random Forest, Linear SVM) not improving at all over the naive model,[19] while others improving to a mediocre degree (Logistic regression, kNN, Decision Tree). However, the *Neural Net* shows a respectable improvement over the naive model, raising accuracy from 0.52 to *0.67*, precision from 0.27 to *0.63* and recall from 0.52 to *0.66*, as well as the F-measure from 0.17 to *0.54*.

It is insightful that the quality goes slightly down when *leaving out tags* among the explanatory features (Neural Net accuracy: *0.62*, F: *0.40*). At the same time, other classifiers become better than neural networks (Naive Bayes accuracy: *0.64*, F: *0.47*). Thus, it seems that tags

[16] http://scikit-learn.org

[17] See Section 3.2. A more sophisticated approach could test different depths using subsumption class hierarchies.

[18] https://pythonhosted.org/lda/

[19] Estimation based on majority class.

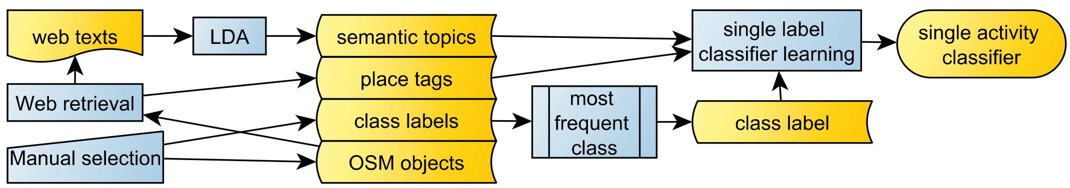

Fig. 11. *slLDA* approach to place affordance modeling. Blue boxes are modeling processes, yellow ones are results.

Table 4
Variants of the slLDA model. We varied the model along each dimension combining parameters.

Classifier learning algorithms	Text sources	Place tags	Semantic depth	Min class size
Logistic regression (C = 1e5), Nearest Neighbors (k = 5), Linear SVM (C = 0.025), RBF SVM (gamma = 2, C = 1), Gaussian Process (1.0 * RBF(1.0), warm start), Decision Tree (max_depth = 5), Random Forest (max_depth = 5, n_est = 10, max_feat = 1), Neural Net (alpha = 1), AdaBoost, Naive Bayes	Scraped web texts, Google reviews	Used, not used	Activity class with referents, without referents	0, 5

Fig. 12. Word clouds for 18 topics from training LDA on web scraped texts.

play an important role as explanatory features for all classifiers to reach highest scores.

A similar picture occurs when using Google reviews instead of web scraped texts (Table 6). Note that due to the more skewed class distribution of these examples, the naive model has already a very high accuracy, which is still overcome by the Neural Net with *0.78*. The *fitting quality* of the neural net is high over all classes, as can be seen by its confusion matrix for estimating activity classes on 136 web scraped texts (Table 7).

When increasing the semantic depth of labels by including 61 activity/referent combinations (such as *ulo:Eating*‖*ulo:Pastry*), accuracy of the Neural Net went down to *0.27* (F score of *0.22*), which only slightly improves over the naive model (accuracy *0.20*, F *0.07*). It seems that the amount and semantic coverage of texts and tags simply does not yield enough information for this purpose.

5.2. mlLDA

For mlLDA *ordinary multi-label classification* approaches were used (Madjarov, Kocev, Gjorgjevikj, & Džeroski, 2012) in a pipeline with *unsupervised LDA* (Fig. 13). In contrast to slLDA, this allowed us to use multiple labels at the same time and has the advantage that we can use non-text features such as tags as explanatory variables in addition to latent topics. Still, as in the first approach, topics irrelevant for classification might occur. In multi-label classification, each instance comes with a list of class labels given as an indicator matrix.[20] There are two basic learning strategies: (1) either single-label classifier algorithms can be adapted specifically to the multi-label problem. A common example is "Multi-label kNN (MLKNN)", which is an adaption of kNN (Zhang & Zhou, 2007). Or (2) one can *transform* the problem into multiple single-label classification problems (Madjarov et al., 2012).

5.2.1. Model parametrization and feature selection

In our experiment, we tested the algorithmic variants as given in Table 8, where *MLKNN* is an algorithmic adaption, while all other classifiers are single-label wrapped into a *label power set* using the *scikit-multilearn* package.[21] The latter turned out to be the most successful transformation method. The LDA method was equivalent to the previous version as well as all other parameter variations, except that we did not restrict class sizes.

[20] A matrix of instances against classes, where each cell is 1 in case the instance is labeled with the class

[21] http://scikit.ml/api/index.html

Table 5
Classification quality of slLDA activity class labeling, using web texts with OSM/Google tags and classes ≥ 5.

	naive model	logistic regr.	kNN	Linear SVM	RBF SVM	Gauss. proc.	Decision Tree	Random Forest	Neural Net	Ada Boost	Naive Bayes
Accuracy	0.52	0.60	0.58	0.52	0.56	0.6	0.58	0.54	**0.67**	0.56	0.49
Std.dev.	0.04	0.13	0.13	0.04	0.08	0.1	0.07	0.1	0.13	0.11	0.16
W.prec.	0.27	0.64	0.61	0.27	0.51	0.46	0.56	0.38	**0.63**	0.58	0.62
W.recall	0.52	0.61	0.58	0.52	0.56	0.6	0.6	0.57	**0.66**	0.56	0.49
F	0.17	0.49	0.48	0.17	0.39	0.32	0.45	0.21	**0.54**	0.48	0.47

Bold indicates the best performing classification

Table 6
Classification quality of slLDA activity class labeling, using Google review texts with OSM/Google tags and classes ≥ 5.

	Naive model	Logistic regr.	kNN	Linear SVM	RBF SVM	Gauss. proc.	Decision Tree	Random Forest	Neural Net	Ada Boost	Naive Bayes
Accuracy	0.67	0.72	0.77	0.67	0.67	0.67	0.64	0.69	**0.78**	0.65	0.62
Std.dev.	0.01	0.17	0.13	0.01	0.01	0.01	0.59	0.05	0.17	0.12	0.22
W.prec.	0.45	0.68	0.7	0.45	0.45	0.45	0.59	0.45	**0.66**	0.55	0.64
W.recall	0.67	0.72	0.77	0.67	0.67	0.67	0.66	0.69	**0.77**	0.65	0.62
F	0.27	0.54	0.59	0.27	0.27	0.27	0.66	0.3	**0.56**	0.4	0.5

Bold indicates the best performing classification

Table 7
Confusion matrix of Neural Net classification using web scraped texts and tags.

		Predicted labels				
		Shopping	Eating	Watching	Drinking	Total
Actual labels	Shopping	**28**	3	0	2	33
	Eating	3	**66**	0	1	70
	Watching	2	1	**12**	0	15
	Drinking	2	7	0	**9**	18
	Total	35	77	12	12	136

Bold indicates the best performing classification

Table 8
Variants of the mlLDA model. We varied the model along each dimension combining parameters.

Classifier learning algorithms	Text sources	Place tags	Semantic depth
Logistic Regression (C = 1e5), MLKNN (k = 5, s = 1.0), Decision Tree (max_depth = 5), Extra Tree (max_depth = 5), Nearest Neighbors (k = 10), Neural Net (alpha = 1), Random Forest (max_depth = 5, n_est = 10, max_feat = 1), Naive Bayes, RBF SVM (gamma = 2, C = 1), Linear SVM (C = 0.025)	Scraped web texts, Google reviews	Used, not used	Activity class with referents, without referents

5.2.2. Results and evaluation

To measure the classification quality, we performed *10-fold cross validation* and used the *weighted* variants of the standard quality measures (accuracy, weighted precision, weighted recall, and weighted F1). For these measures, estimation qualities for each class are weighted by the size of that class. Standard multi-label quality measures take into account label order, however, which is not relevant in our case.[22] Thus, we also calculated order-neutral measures including 1) the *coverage error*, which computes the average number of labels that have to be included in the final prediction such that all true labels are covered, and 2) the *Jaccard similarity score*, which computes the average Jaccard similarity coefficient over all instances. The latter coefficient is the ratio of the number of shared (true and predicted) labels, divided by the size of the union of the true and the predicted labels. When looking at the performance of the models on the Web scraped texts including tags (Table 9), we see that again the Neural Net outperforms the other algorithms and raises the accuracy of the naive model from *0.24* to *0.4*. The more tolerant coverage error drops from almost *14* to *10.7*, and the Jaccard score is raised from *0.34* to over *0.52*. This means that *on average*, more than half of all the class labels present are correctly predicted by the model. An only slightly worse quality is reached by kNN, the naive Bayes classifier and logistic regression. Surprisingly, the adapted MLKNN algorithm did not particularly stand out.

Leaving out tags considerably deteriorated models across all classifiers. Most classifiers were not capable anymore to raise the level beyond a naive guess. The Neural Net accuracy drops to *0.37*, and its Jaccard score to *0.45*.

We also tested the approach on *Google review* texts. In this case, the naive accuracy of *0.31* was bested by the Neural Network *0.38*, and the Jaccard score improved from *0.44* to *0.59*. It seems, however, that the sample size in this case is too small for the task. Finally, we tested with *highest semantic resolution* (including referent classes). As expected, this turned down the quality considerably. The Neural Net was able to raise the accuracy in this case from *0.05* only to *0.11*, and the Jaccard score from *0.06* to *0.15*.

5.3. LLDA

LLDA is a variant of LDA that uses the additional information of user-contributed labels to identify topics from texts in a supervised manner, so that there is a one-to-one correspondence between labels and topics (Ramage, Hall, Nallapati, & Manning, 2009). Since a text can have more than one label, it is modeled as a mixture over those corresponding topics. This has the advantage that topics are chosen in a way that optimizes discrimination between class labels. The drawback lies in the fact that LLDA is restricted to texts as features. This approach is described in Fig. 14.

5.3.1. Model parametrization and feature selection

Similar to the LDA model, we restricted models to those classes with at least 5 instances. For parameterization of the *LLDA algorithm*, the α

[22] Since the order of activities can be disregarded.

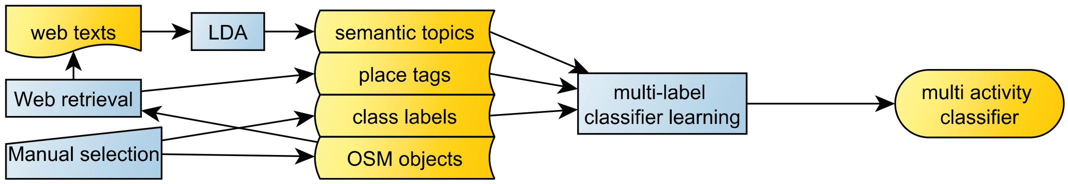

Fig. 13. *mlLDA* approach to place affordance modeling.

Table 9
Multi-label classification quality of activity classes, using web scraped texts with OSM/Google tags.

	Naive model	Logistic regr.	MLKNN	Decision Tree	Extra Tree	KNN	Neural Net	Random Forest	Naive Bayes	RBF SVM	Linear SVM
Accuracy	0.24	0.33	0.31	0.22	0.23	0.36	**0.40**	0.30	0.31	0.31	0.24
Std.dev.	0.04	0.04	0.04	0.02	0.04	0.02	0.03	0.06	0.06	0.05	0.05
W.prec.	0.14	0.53	0.57	0.45	0.33	0.48	**0.58**	0.36	0.6	0.37	0.15
W.recall	0.31	0.5	0.43	0.44	0.34	0.43	**0.48**	0.38	0.51	0.36	0.31
F	0.17	0.48	0.47	0.42	0.31	0.43	**0.49**	0.33	0.52	0.33	0.2
Coverage	13.9	10.56	11.95	11.11	12.94	10.88	**10.71**	13.01	10.68	12.47	13.9
Jaccard	0.34	0.47	0.44	0.41	0.36	0.47	**0.52**	0.38	0.47	0.41	0.34

Bold indicates the best performing classification

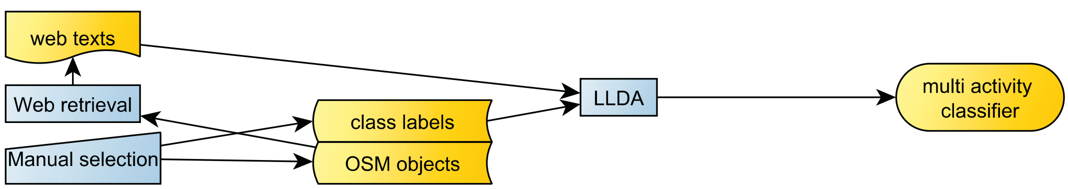

Fig. 14. *LLDA* approach to place affordance modeling.

Table 10
Top words for LLDA topics trained using general activity class labels.

Eating	Shopping	Drinking	Watching	Sitting	Listening	Dancing
eten	onze	koffie	langhuis	augustus	dienst	club
sfeer	wij	onze	stichting	kerk	live	bommel
service	winkel	we	jaar	mizu	podia	jack
kwaliteit	augustus	bier	path	zoeken	uur	underground
Prijs	we	caf	kunst	zwolse	thor	dj
Drinken	cookies	wij	pand	kunt	verkocht	wils
We	nederland	plek	nieuwe	indebuurt	gesteld	evenementen
Schaal	website	barista	kunstenaars	zon	diverse	reactie
restaurant	waar	kunt	galerie	terras	show	bloopers

Table 11
LLDA classification quality of activity classes, using Web scraped texts.

	Naive model	LLDA
Accuracy	0.24	0.34
W.prec.	0.14	0.44
W.recall	0.31	0.45
F	0.17	0.42
coverage	13.9	10.5
jaccard	0.34	0.42

hyperparameter, used to determine the distribution of topics over documents, was set to 0.001 and training was performed with 1000 iterations.

5.3.2. Results and evaluation

The LLDA model was trained on the same Web corpus that was described in Section 5.1.2 and Table 9. As shown in Table 10 many top words for topics that were trained based on general activity labels correspond with the labeled activities. For example, the Dutch words "eten" (eating), "sfeer" (atmosphere) and "schaal" (plate) correspond to *Eating* in a restaurant, "zon" and "terras" correspond to *Sitting* on a terrace in the sun, and "live", "podia" and "show" correspond to *Listening* to concerts. Words include also sensible local venues names, like "Mizu" (a cocktail bar) and "Thor"(a music club).

In order to evaluate how well labels are correctly applied to texts, we performed a *10-fold cross validation* on the data. Since the LLDA prediction results in probabilities for each label rather than label indicators (i.e., true or false for each label), an additional threshold value of 0.2 was applied to turn LLDA into a multi-label classifier comparable to Table 9. Table 11 shows these results for LLDA.

Based on this method the quality of LLDA is better than the naive model but is not as high as the best performing multi-label classifiers shown in Section 5.2.2.

6. Discussion, outlook and conclusion

The results of our study add to more than one direction of research on place and urban leisure space. We therefore discuss them separately for the three questions posed in the introduction.

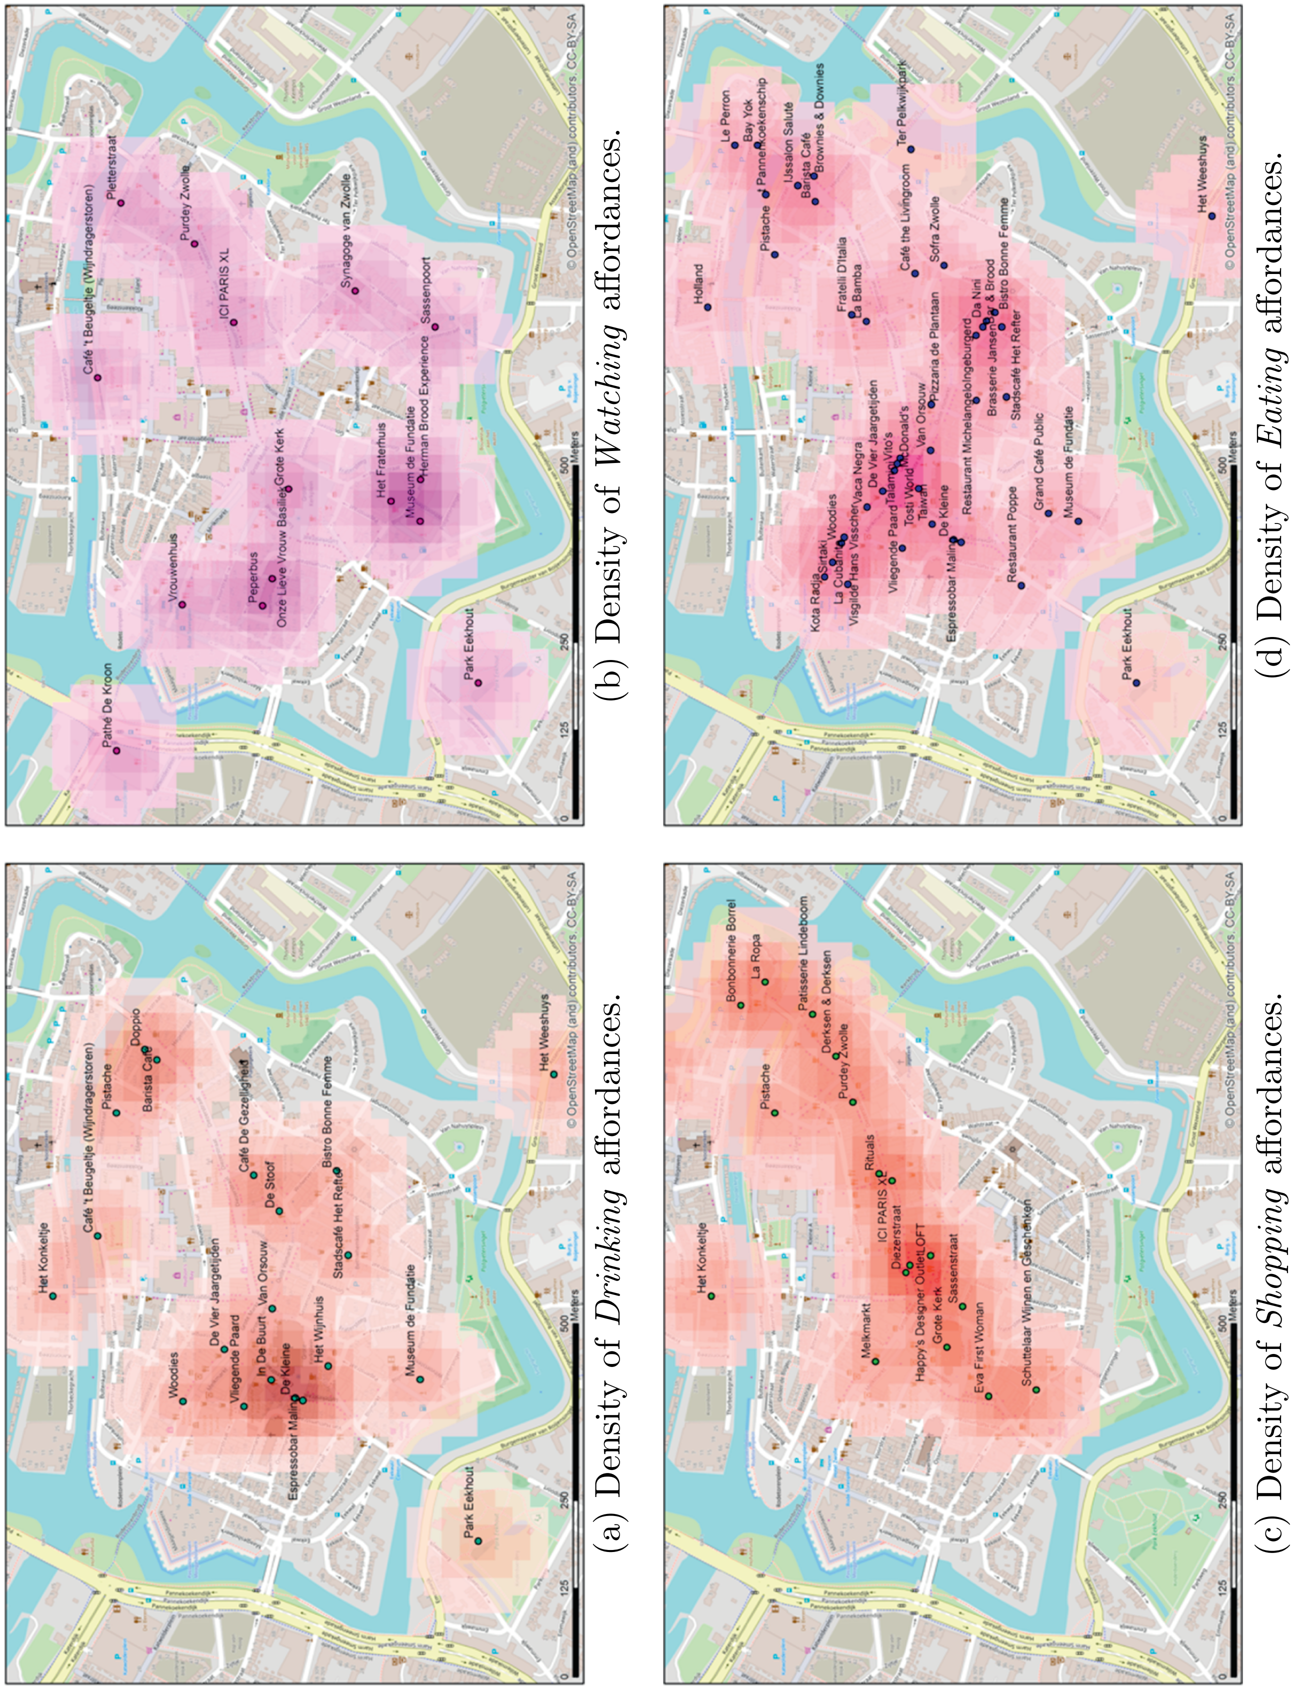

(a) Density of *Drinking* affordances. (b) Density of *Watching* affordances. (c) Density of *Shopping* affordances. (d) Density of *Eating* affordances.

Fig. 15. Kernel density maps of place affordances in the city center of Zwolle.

6.1. Quality of extracting leisure activity potentials from Web sources

The results show that even though place affordance estimation down to the semantic level of activity referents (what kind of food is eaten) is a hard problem, a particular machine learning approach (neural nets trained on semantic topics from websites and review texts and making use of place tags) is able to achieve a considerable quality on the level of activity classes. The most important activity class can be estimated with an accuracy of ~ 0.8 (using review texts) and ~ 0.7 (using webtexts), where accuracy, precision, and recall are significantly better than a naive model. Also the more challenging multi-label problem can be handled in this way, where on average more than half of all true activity class labels are correctly suggested by our model. Surprisingly, supervised multi-label models specifically designed for the task, such as MLKNN and LLDA, did not perform any better. Social media tags from OSM and Google places were an indispensable source for the task.

In summary, the variety of place affordances and functions can be estimated with a moderate to good quality by combining various texts and tag sources from the web. However, this can be done only on a rather coarse semantic level of activity classes. This demonstrates that there remains indeed a large gap between the alleged benefits of readily available data on the Web and the difficulties involved in making scalable and reliable statements about affordances and experiential qualities that urban places actually offer, c.f. Section 2.2 and 2.3.

To extend our model to other cities, our study should be considerably extended. Since our manually labeled sample was rather restricted in size and based on a single city, many meaningless topics based on text artifacts (such as web formatting) remained, which might disappear with larger samples. Furthermore, a lot of missing items were produced by Web scraping due to service quality and incompleteness of social media data. A larger issue related to this is how a sufficient text quality of web sources can be assured when expanding the method beyond our manually curated sample. For this purpose, studies on principle Web data quality are needed (Ostermann & Granell, 2017), in order to compare different sources, and in order to assess their sample biases (Caliskan et al., 2017). Future work could build on our results in order to realize a repository of urban activity models with labeled place data across cities. It would then become possible to compare the leisure potentials across these cities.

6.2. Studying the landscape of leisure activity potentials in Zwolle

Our approach allows us to translate urban texts into a technical machine-readable vocabulary, which can be used in GIS analysis to reveal different spatial patterns of specialization and mixing of place affordances. For example, in Fig. 15 we have produced high resolution kernel density maps about four different place affordances in the city centre of Zwolle, namely drinking, eating, shopping, and watching.

On these affordance maps, one can discover that places for eating and places for shopping are concentrated around different crossing streets and partly overlap in the old town. Places combining the affordance of eating and drinking are mostly concentrated on the square surrounding the Sint-Michaëls church. When visitors do a Google maps search of Zwolle and zoom in on its city centre, they may discover this particular spatial pattern of specialization in terms of tags. However, they will have great difficulties in discovering side activities and concealed place qualities. Google places category tags hide that some places afford multiple kinds of activities. For example, as can be seen in Fig. 15, the Museum de Fundatie and the Sint-Michaëls church provide not only the obvious watching affordances—e.g. in terms of architectural features of the buildings and the visual arts inside. The museum is also depicted as a place for drinking and eating—due to a small bar on its third floor—and the church as a place for shopping—due to a book market inside to financially support its maintenance.

Furthermore, less prominent places are also reliably captured. In fact, the inner city of Zwolle is covered almost completely in our model. Fig. 15 does not only display the main tourist highlights, but also architectural highlights "off the beaten track". For example, chain store shops like 'Ici Paris XL' and 'Purdey' are discovered as interesting places for watching because they are both housed in buildings of historical value. The 'Ici Paris XL' store is housed in a 17th century building with a Classicist style and the 'Purdey' store in a 20th century building with a Jugendstil style. At the same time, the Eekhout park provides various monuments, an open air terrace and musical performances, which seem known and quite popular among locals but probably not among visitors yet.

This demonstrates that in contrast to Google maps or OSM, our affordance maps display a large diversity of leisure affordances in the city centre of Zwolle which are not captured by place names or tags.

6.3. Coping with the long tail and soft spot of urban research

Our affordance maps of Zwolle (Fig. 15) demonstrate the potential to account for several challenges in existing data science approaches to urban leisure and tourism studies and for related city policies and planning.

First, it may be used to avoid the quantitative bias in the currently available geographic information of social media data (ratings, clicks, number of reviews), providing a platform to dig into more "qualitative" soft and subtle information sources, such as perceptions and experiences of place affordances. Texts provide a much richer semantic resource for this purpose compared to only tags.

Second, it provides a way to avoid the bias coming from overlooking the *long tail* in place popularity and place categorization. Place popularity reinforces attention and visitor flows to the most popular sites (Scott & Orlikowski, 2012; van der Zee et al., 2018), while place categorization reduces urban space to the most obvious things to do there.

Third, in combining different sources within distinct features, we avoid the bias coming from single sources, such as a single social media site or first-hand website. This provides a more balanced image as different types of sources are known to present different aspects of places (Xiang & Gretzel, 2010).

This provides fascinating opportunities for future research in urban leisure and tourism and for innovating city policies and planning. For example, it helps overcoming a one-size-fits-all approach to city marketing. Extracting information of high spatio-temporal and semantic resolution, may support more personalized urban exploration of leisure seekers. It may also show a way to better direct the crowd of visitors and residents in order to distribute them over various places within the city (compare Section 2). Furthermore, our type of urban spatial analysis can enrich tracking and quantitative spatial behavior studies with precious qualitative information, which is currently not available or needs to be obtained via costly surveys and interviews.

However, our machine learning study also showed that the semantic resolution and accuracy has limits imposed by the quality of sample texts. For this reason, developing purely data-driven strategies for urban planning should be approached with care. It might be the case that information about needed place affordances in a city are just not explicit in the text sources or that they are not covered in detail. Future research should therefore measure the quality of existing Web data sources for a given semantic level. For example, while our results illustrate the gastronomic appeal of Zwolle, it is unclear whether our method can also extract the availability of vegetarian food.

References

Adams, B., & Janowicz, K. (2015). Thematic signatures for cleansing and enriching place-related linked data. *International Journal of Geographical Information Science, 29*, 556–579.

Adams, B., & McKenzie, G. (2013). Inferring thematic places from spatially referenced natural language descriptions. *Crowdsourcing geographic knowledge* (pp. 201–221). Springer.

Alazzawi, A. N., Abdelmoty, A. I., & Jones, C. B. (2012). What can I do there? Towards the automatic discovery of place-related services and activities. *International Journal of Geographical Information Science, 26*, 345–364.
Ashworth, G., & Page, S. J. (2011). Urban tourism research: Recent progress and current paradoxes. *Tourism Management, 32*, 1–15.
Ballatore, A. (2014). The search for places as emergent aggregates. *Spatial Search-Specialist meeting*. Santa Barbara Santa Barbara (CA), USA: Center for Spatial Studies, University of California.
Ballatore, A., Wilson, D. C., & Bertolotto, M. (2013). *A survey of volunteered open geo-knowledge bases in the semantic web, in: Quality issues in the management of web information.* Springer93–120.
Barrett, H., & Hall, T. (2018). *Urban Geography.* Routledge.
Batty, M. (2013). Big data, smart cities and city planning. *Dialogues in Human Geography, 3*, 274–279.
Blei, D. M. (2012). Probabilistic topic models. *Communications of the ACM, 55*, 77–84.
Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of Machine Learning Research, 3*, 993–1022.
Buhalis, D., & Amaranggana, A. (2013). *Smart tourism destinations, in: Information and communication technologies in tourism.* Springer553–564.
Bureau voor Economische Argumentatie, & Zwolle bruist - De strategische agenda binnenstad (2017). *2017–2022,*. Bureau voor Economische Argumentatie.
Caliskan, A., Bryson, J. J., & Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science, 356*, 183–186.
Cataldi, M., Ballatore, A., Tiddi, I., & Aufaure, M.-A. (2013). Good location, terrible food: Detecting feature sentiment in user-generated reviews. *Social Network Analysis and Mining, 3*, 1149–1163.
Chesnokova, O., & Purves, R. S. (2018). From image descriptions to perceived sounds and sources in landscape: Analyzing aural experience through text. *Applied Geography, 93*, 103–111.
Clon, M. (2013). *Met het oog op de tijd.* Sociaal en Cultureel Planbureau.
Crosbie, T. (2006). Using activity diaries: Some methodological lessons. *Journal of Research Practice, 2*, 1.
Fabrikant, S. I., & Buttenfield, B. P. (2001). Formalizing semantic spaces for information access. *Annals of the Association of American Geographers, 91*, 263–280.
Farrahi, K., & Gatica-Perez, D. (2011). Discovering routines from large-scale human locations using probabilistic topic models. *ACM Transactions on Intelligent Systems and Technology (TIST), 2*, 3.
Friedman, J., Hastie, T., & Tibshirani, R. (2001). The elements of statistical learning. *Springer series in statistics, New York*. *Vol. 1*.
Gangemi, A. (2013). A comparison of knowledge extraction tools for the semantic web. *Extended Semantic Web Conference* (pp. 351–366). Springer.
Gangemi, A., & Presutti, V. (2009). Ontology design patterns. *Handbook on ontologies* (pp. 221–243). Springer.
Girardin, F., Calabrese, F., Dal Fiore, F., Ratti, C., & Blat, J. (2008). Digital footprinting: Uncovering tourists with user-generated content. *IEEE Pervasive Computing, 7*.
Goodchild, M. F. (2007). Citizens as sensors: The world of volunteered geography. *GeoJournal, 69*, 211–221.
Gospodini, A. (2006). Portraying, classifying and understanding the emerging landscapes in the post-industrial city. *Cities, 23*, 311–330.
Guo, Y., Barnes, S. J., & Jia, Q. (2017). Mining meaning from online ratings and reviews: Tourist satisfaction analysis using latent dirichlet allocation. *Tourism Management, 59*, 467–483.
Hanna, S. A., & Rowley, J. (2015). Rethinking strategic place branding in the digital age. *Rethinking place branding* (pp. 85–100). Springer.
Hannigan, J. (1998). *Fantasy city: Pleasure and profit in the postmodern metropolis.* Routledge.
Hasan, S., & Ukkusuri, S. V. (2014). Urban activity pattern classification using topic models from online geo-location data. *Transportation Research Part C, 44*, 363–381.
Hollenstein, L., & Purves, R. (2010). Exploring place through user-generated content: Using flickr tags to describe city cores. *Journal of Spatial Information Science, 2010*, 21–48.
Hu, Y., Gao, S., Janowicz, K., Yu, B., Li, W., & Prasad, S. (2015). Extracting and understanding urban areas of interest using geotagged photos. *Computers, Environment and Urban Systems, 54*, 240–254.
Hudson, S., & Thal, K. (2013). The impact of social media on the consumer decision process: Implications for tourism marketing. *Journal of Travel & Tourism Marketing, 30*, 156–160.
Jonietz, D. (2016). *From space to place: A computational model of functional place.* Ph.D. thesisUniversity of Augsburg.
Jordan, T., Raubal, M., Gartrell, B., & Egenhofer, M. (1998). An affordance-based model of place in gis. *8th Int. Symposium on Spatial Data Handling, SDH*. *98*. *8th Int. Symposium on Spatial Data Handling, SDH* (pp. 109–). .
Kitchin, R. (2014). The real-time city? Big data and smart urbanism. *GeoJournal, 79*, 1–14.
Kwan, M.-P. (2002). Time, information technologies, and the geographies of everyday life. *Urban Geography, 23*, 471–482.
Lorentzen, A. (2009). Cities in the experience economy. *European Planning Studies, 17*, 829–845.
Lu, W., & Stepchenkova, S. (2015). User-generated content as a research mode in tourism and hospitality applications: Topics, methods, and software. *Journal of Hospitality Marketing & Management, 24*, 119–154.
Madjarov, G., Kocev, D., Gjorgjevikj, D., & Džeroski, S. (2012). An extensive experimental comparison of methods for multi-label learning. *Pattern Recognition, 45*, 3084–3104.
Marine-Roig, E., & Clavé, S. A. (2015). Tourism analytics with massive user-generated content: A case study of Barcelona. *Journal of Destination Marketing & Management, 4*, 162–172.
McKenzie, G., Janowicz, K., Gao, S., Yang, J.-A., & Hu, Y. (2015). Poi pulse: A multi-granular, semantic signature–based information observatory for the interactive visualization of big geosocial data. *Cartographica, 50*, 71–85.
Montello, D. R., Goodchild, M. F., Gottsegen, J., & Fohl, P. (2003). Where's downtown?: Behavioral methods for determining referents of vague spatial queries. *Spatial Cognition & Computation, 3*, 185–204.
NBTC-NIPO Research (2016). *Continu VrijeTijdsOnderzoek,*. NBTC-NIPO Research.
Ostermann, F. O., & Granell, C. (2017). Advancing science with VGI: Reproducibility and replicability of recent studies using VGI. *Transactions in GIS, 21*, 224–237.
Purves, R., Edwardes, A., & Wood, J. (2011). Describing place through user generated content. *First Monday, 16*.
Purves, R. S., & Derungs, C. (2015). From space to place: Place-based explorations of text. *International Journal of Humanities and Arts Computing, 9*, 74–94.
Ramage, D., Hall, D., Nallapati, R., & Manning, C. D. (2009). Labeled lda: a supervised topic model for credit attribution in multi-labeled corpora. *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*. *Volume 1*. *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing* (pp. 248–256). Association for Computational Linguistics.
Scheider, S., & Janowicz, K. (2014). Place reference systems. *Applied Ontology, 9*, 97–127.
Scott, S. V., & Orlikowski, W. J. (2012). Reconfiguring relations of accountability: Materialization of social media in the travel sector. *Accounting, Organizations and Society, 37*, 26–40.
Spierings, B. (2006). *Cities, consumption and competition: The image of consumerism and the making of city centres.* Ph.D. thesisNijmegen: Radboud University.
Spierings, B. (2013). Fixing missing links in shopping routes: Reflections on intra-urban borders and city centre redevelopment in Nijmegen, the Netherlands. *Cities, 34*, 44–51.
Szalai, A. (1972). *The use of time: Daily activities of urban and suburban populations in 12 countries.* Mouton.
Townsend, A. M. (2013). *Smart cities: Big data, civic hackers, and the quest for a new utopia.* WW Norton & Company.
Xiang, Z., & Gretzel, U. (2010). Role of social media in online travel information search. *Tourism Management, 31*, 179–188.
Ye, Q., Zhang, Z., & Law, R. (2009). Sentiment classification of online reviews to travel destinations by supervised machine learning approaches. *Expert Systems with Applications, 36*, 6527–6535.
van der Zee, E., Bertocchi, D., & Vanneste, D. (2018). Distribution of tourists within urban heritage destinations: A hot spot/cold spot analysis of tripadvisor data as support for destination management. *Current Issues in Tourism,* 1–22.
van der Zee, E., van der Borg, J., & Vanneste, D. (2017). The destination triangle: Toward relational management. *Knowledge transfer to and within Tourism: Academic, Industry and Government Bridges, Emerald Publishing Limited* (pp. 167–188). .
Zhang, M.-L., & Zhou, Z.-H. (2007). Ml-knn: A lazy learning approach to multi-label learning. *Pattern Recognition, 40*, 2038–2048.
Zukin, S. (1998). Urban lifestyles: Diversity and standardisation in spaces of consumption. *Urban Studies, 35*, 825–839.