

POINTREC: A Test Collection for Narrative-driven Point of Interest Recommendation

Jafar Afzali
University of Stavanger
j.afzali@stud.uis.no

Aleksander Mark Drzewiecki
University of Stavanger
am.drzewiecki@stud.uis.no

Krisztian Balog
University of Stavanger
krisztian.balog@uis.no

ABSTRACT

This paper presents a test collection for contextual point of interest (POI) recommendation in a narrative-driven scenario. There, user history is not available, instead, user requests are described in natural language. The requests in our collection are manually collected from social sharing websites, and are annotated with various types of metadata, including location, categories, constraints, and example POIs. These requests are to be resolved from a dataset of POIs, which are collected from a popular online directory, and are further linked to a geographical knowledge base and enriched with relevant web snippets. Graded relevance assessments are collected using crowdsourcing, by pooling both manual and automatic recommendations, where the latter serve as baselines for future performance comparison. This resource supports the development of novel approaches for end-to-end POI recommendation as well as for specific semantic annotation tasks on natural language requests.

CCS CONCEPTS

• **Information systems** → **Environment-specific retrieval**; *Users and interactive retrieval*; *Recommender systems*;

KEYWORDS

POI recommendation; point of interest; cold start recommendation; contextual suggestions; narrative-driven recommendation

ACM Reference Format:

Jafar Afzali, Aleksander Mark Drzewiecki, and Krisztian Balog. 2021. POINTREC: A Test Collection for Narrative-driven Point of Interest Recommendation. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '21)*, July 11–15, 2021, Virtual Event, Canada. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3404835.3463243>

1 INTRODUCTION

Point of interest (POI) recommendation entails the task of recommending locations (e.g., restaurants, arts, and entertainment) to users based on their interest profiles. It is a particular item recommendation task, which has attracted considerable interest over the past decade. POI recommendation has especially gained popularity with the emergence of location-based social networks (LBSNs), like

Yelp, Gowalla, or Foursquare, and is seen as an essential feature of those platforms [6, 7, 25]. On LBSNs, users share their social experience via check-ins and ratings/reviews. This data *implicitly* captures their preferences on locations and provides the basis for future recommendations. Crucially, these general preferences are often insufficient when making recommendations; context may introduce additional constraints, not necessarily inclined to personal preferences [1]. For example, a person may be a frequent visitor to nightlife spots, but would prefer very different type of places during a weekend family getaway.

Aiding POI recommendation by incorporating contextual information was the very focus of the Contextual Suggestions track at the Text Retrieval Conference (TREC) [9, 12]. There, systems are required to make suggestions for a particular person in a specific context (e.g., a city and trip type) based on user profiles, which comprise stated preferences on POIs in different cities and contexts. The particular needs and preferences of users are supposed to be *inferred* from their interests. This, however, appears to be a challenge, and “not many successful participants took into account context in their proposed approaches” [1]. Our objective with this work is to eliminate difficulties associated with the extraction of contextual preferences and constraints, by making these more explicit, thereby facilitating research on novel contextual POI recommendation methods.

This paper addresses a POI recommendation scenario, where a contextual information need is stated *explicitly* in natural language. This scenario has been termed as *narrative-driven recommendation* in [4]. However, unlike in [4], past user transactions are not available in our setting (nor is user profile information, apart from what is explicitly stated in the description). This renders it a particularly difficult and interesting problem. It may be argued that this is more of a *search* task than a recommendation one, since there is an explicit expression of the user need as opposed to implicit preferences captured by user history. At the same time, the user is asking for *recommendations*, i.e., they are seeking for suggestions to choose from. Their information needs are too complex to be expressed as search queries that would be understood by search engines. Rather, these are exactly the types of requests that one would make to a conversational assistant. One particularly challenging aspect of this problem is natural language understanding, i.e., extracting and interpreting preferences that are expressed in free text form. Our work facilitates research in this area by providing semantically rich ground truth annotations of user preferences.

We present the POINTREC collection consisting of the following:

- A collection of 695K POIs, spanning over 23 countries and over 2K cities, which are collected from a popular online directory (Yelp), and are further linked to a geographical knowledge base (GeoNames) and enriched with relevant web snippets.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGIR '21, July 11–15, 2021, Virtual Event, Canada

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8037-9/21/07...\$15.00

<https://doi.org/10.1145/3404835.3463243>

- A set of 112 information needs originating from real users, each with a text description detailing the context and preferences of the user. These natural language descriptions are enriched with metadata annotations and, if possible, linked to structured knowledge repositories (including category, group, location information, constraints, and examples).
- Graded relevance assessments collected via crowdsourcing.

It is worth pointing out that collaborative filtering techniques, which have been the predominant approach to POI recommendation to date [6, 13–15, 25], are not applicable in our scenario, given the lack of user-item interactions. Instead, we approach the task using a retrieval-based approach, where constraints are expressed as structured queries on fielded representations of items. We offer a starting point by implementing and comparing the performance of three variants of this baseline.

The POINTREC collection and our baselines are made available at <https://github.com/iai-group/sigir2021-pointrec>.

2 RELATED WORK

The emergence of location-based social networks (LBSNs) brought about the need for POI recommendation as an essential service, in order to help users discover and explore new places [6, 7, 25]. POI recommendation in LBSNs “distinguishes itself from traditional item recommendation [...] via geographical influence among POIs” [21]. While traditional collaborative filtering approaches offer a natural starting point, these need to be extended to incorporate spatial, temporal, and social information [10, 22, 25, 27]. Another major challenge is data sparsity. While there are millions of POIs, each user visits only a handful of them, making the user-POI interaction matrix extremely sparse [26]. Latent factor models, and in particular matrix factorization approaches, are often adopted to alleviate data sparsity [6, 13–15, 17]. Alternatively, probabilistic generative models may also be employed [26]. Most recent work has shown the effectiveness of embedding-based approaches for POI recommendation [18, 22, 23]. Evaluation datasets are created by crawling LBSN data via public APIs; these crawls are often shared by individual researchers, e.g., for Gowalla [6, 8, 16] and Foursquare [11, 24, 27]. These datasets are often restricted to specific geographic regions.

A key characteristic that makes POI recommendation a fertile ground for research is the inherent need for contextualization. Context may be a previous POI (or category of POI) [3, 15], temporal state (e.g., work vs. leisure time) [17, 27], travel time [5], or weather [5]. Recommendations may even be shown proactively to users when the inferred context seems appropriate [3, 5]. The contextual appropriateness of locations may be estimated using heuristics [5] or cast as a binary classification problem [1].

The TREC Contextual Suggestions (TREC CS) track [9, 12] provides a reusable test collection by considering five contextual dimensions: city (e.g., Boston), trip type (e.g., holiday), duration (e.g., weekend trip), type of group (e.g., friends), and a season the trip will occur in (e.g., summer). The POI collection covers attractions from 272 North American cities. POIs are represented by their URL, title, and city ID. There are several key differences between TREC-CS and our POINTREC test collection:

- The information needs at TREC CS are “synthetic” in the sense that they are generated by crowdsourcing using an imaginary scenario. Ours, on the other hand, are real information needs by real people with an actual intent to travel.
- The contexts of information needs as well as user preferences in our collection are described in natural language in detail, and have been manually annotated along several dimensions; in case of TREC CS, user preferences need to be inferred. On the other hand, our collection does not include past user history (unless references are made to prior experiences in the description), thereby making it a cold start problem.
- Our POI collection covers 23 countries, while TREC CS deals with North American cities only.

3 INFORMATION NEEDS

POINTREC comprises 112 information needs, spanning over four main categories (Section 3.1). These represent actual information needs of different people, which have been collected from community question answering and forum sites (Section 3.2). Each of these requests has been annotated with a rich set of metadata attributes (Section 3.3). Both these steps have been performed manually by the authors of the paper to ensure high data quality. The resulting collection exhibits a diverse set of needs and attributes (Section 3.4).

3.1 Categories

We collected information needs in four main categories: *Active life*, *Arts and entertainment*, *Restaurants and food*, and *Nightlife*. This particular selection was motivated by looking at the most popular top-level categories in Yelp.¹

Within these main categories, a number of sub-categories are distinguished, following Yelp’s category system.² Due to inappropriateness, a few sub-categories targeting adult entertainment have been actively avoided. A sample from the list of subcategories is given in Table 1.

3.2 Collecting Information Needs

Information needs were collected from two main sources: Yahoo! Answers (24 in total) and Reddit (88 in total).

In Yahoo! Answers, potential posts were identified by searching for recommendations specifically for each of the main categories (e.g., by issuing the query “restaurant recommendations”). A similar procedure was used for Reddit, but the search was restricted to specific subreddits that either correspond to a given location (e.g., “Oslo”) or to a specific activity (e.g., “hiking”). A list of relevant subreddits was manually identified.³

3.3 Annotations and Metadata Extraction

A unique property of the POINTREC collection is that information needs are manually annotated with rich metadata. In the beginning, “raw” information needs contain the following fields:

- **URL:** Post where the information need originates from.
- **Title:** The title of the post.
- **Description:** The body of the post, describing the information need.

¹Actually, *Restaurants* and *Food* are two separate categories in Yelp, which have been merged into a single one here.

²https://www.yelp.com/developers/documentation/v3/all_category_list

³<https://github.com/iai-group/sigir2021-pointrec/blob/master/Reddit.md>

Table 1: Category system, with a sample of subcategories for each main category.

Main category	Subcategories
Active life	ATV Rentals/Tours, Airsoft, Amateur Sports Teams, Amusement Parks, Aquariums, Archery, Beaches, Bowling, Fitness & Instruction, Golf, Hiking, Parks, Skydiving, Zoos, ...
Arts and entertainment	Arcades, Art Galleries, Botanical Gardens, Casinos, Castles, Museums, Music Venues, LAN Centers, Performing Arts, Social Clubs, Street Art, Observatories, Wineries, Virtual Reality Centers, ...
Restaurants and food	Coffee & Tea, Desserts, Food Trucks, Internet Cafes, Specialty Food, Wineries, Afghan, African, American (Traditional), Asian Fusion, Beer Garden, Cafes, Comfort Food, Polish, ...
Nightlife	Bar Crawl, Bars, Beer Gardens, Coffeeshops, Music Venues, Dance Clubs, Club Crawl, Pubs, Lounges, Hookah Bars, Gay Bars, Dance Restaurants, Pool Halls, Sports Bars, ...

Table 2: Possible qualifiers expressing constraints for information needs.

Qualifier	Weight	Example	
		Constraints_text	Constraints
MUST (have/be)	3	"I want to see museums"	museums
SHOULD (have/be)	2	"Good service is appreciated"	good service
NICE_TO (have/be)	1	"I like museums, monuments, and arts"	[museums, monuments, arts]
NICE_TO_NOT (have/be)	-1	"I don't prefer EDM, but I can live with it."	Electronic dance music
SHOULD_NOT (have/be)	-2	"I'm not a fan of the big touristy spots"	tourist centric
MUST_NOT (have/be)	-3	"I do not want to see museums"	museums

It is important to note that a single post may give rise to multiple information needs. Specifically, it may happen that the user is seeking recommendations in multiple main categories, countries, or cities. In such cases, a number of separate entries are created, which have identical URL, title, and description, but will differ in their metadata fields. This information is also preserved in the unique **ID**, which follows a particular AAAA-BBB-CC format, where:

- AAAA is four running digits to identify the post where the information need originates from;
- BBB is a counter that tracks if the originating post contains multiple information needs;
- CC is used to identify which category this information need belongs to (AL, AE, RF, or NL).

Next, information needs are enriched with the following metadata:

- **Main Category:** This field specifies which main category the information need belongs to. (By definition, each information need belongs to a single main category.) Note that the main category is not necessarily explicitly listed, e.g., someone looking for hiking recommendations may not specify "Active Life." The identification of main- and sub-categories is done manually, by reading the description and then matching it to one or several categories from Yelp.
- **Sub-categories:** An exhaustive list of all sub-categories that are relevant for the information need.
- **Group:** A description of the group traveling. For information needs where this is not specified, the field is left empty. In all other cases, it is a contiguous piece of text extracted from the description (e.g., "a group of 10" or "with my girlfriend").
- **Constraints_text:** A list of user preferences regarding the POI. This field is copied verbatim from the description (i.e.,

each element of the list is a contiguous piece of text in the description).

- **Constraints:** Same as the field above, with the preferences structured into (qualifier, constraint) pairs. Qualifiers come from a fixed set: MUST, SHOULD, NICE_TO, and their negative counterparts. The natural language formulation of the constraint is possibly simplified. For example, "can't cost an arm and leg" implies that they are not looking for anything too expensive, so it is interpreted as "MUST_NOT too expensive." Table 2 lists the possible qualifiers along with examples.
- **Positive examples:** Some information needs include POIs that the user likes, or, is interested in. These POIs are listed here.
- **Negative examples:** Analogous to the Positive examples field, this is a listing of POIs that are explicitly disliked or indicated as not interesting.
- **Request_text:** A specification of what the user is looking for, copied verbatim from the description. It is worth noting that this field may contain multiple contiguous pieces of text from the description.
- **Request:** Indicates the type of place/activity the user is looking for, e.g., "restaurant," "museum," "hiking," or "bars."
- **Country:** The Alpha-2 code⁴ of the country the information need is located in.
- **State:** This field is primarily used for large countries (USA, Canada, and Australia).
- **City:** The name of the city targeted by the information need.⁵

⁴<https://www.iso.org/glossary-for-iso-3166.html>

⁵For some of the information needs, the city name and country name are not present in the title or description. However, as the information needs were posted in forums dedicated to a specific city, this was inferred.

Table 3: Example information need (ID=0020-000-RF).

<i>Raw information need</i>	
URL	https://answers.yahoo.com/question/index;...
Title	Restaurant recommendation in London near Charing Cross?
Description	I want to take my Dad out for a nice meal for a big birthday (really soon). Wondering if anyone can suggest anywhere? He is basically the typical grumpy old man, likes plenty on his plate, no waiters flapping over you (like good service!), no loud music, good ale on tap is a definite bonus and can't cost an arm and leg. It also has to be in walking distance of Charing Cross as he won't use the tube. Pubs or restaurants would be fine as long as they aren't crowded/noisy and meet the above criteria. Any suggestions?
<i>Metadata enrichment</i>	
Group	my father and I
Main Category	Restaurants and Food
Constraints_text	He is basically the typical grumpy old man ...
Constraints	MUST: ["Good service", "Walking distance of Charing Cross"] SHOULD: ["Ale on tap"] MUST_NOT ["Loud music", "Crowded", "Noisy", "too expensive"]
Request_text	Restaurant recommendation in London near Charing Cross?
Request	Restaurant
Country	GB
City	London

Notice that there are two main types of metadata fields. The first group is essentially a mark-up on the original description text, identifying certain types of preferences (group, constraints_text, request_text, positive/negative examples). The second group of fields correspond to semantic annotations, which involve some degree of linguistic normalization (constraints, request) or restricting values to a controlled vocabulary (categories, country, state, city). Both types of metadata could potentially be extracted automatically; this, however, is left for future work.

Table 3 shows an example information need, including the raw need (top) and the corresponding metadata enrichment (bottom).

3.4 Statistics

A breakdown of the information needs across the main categories is shown in Table 4. It can be seen that information needs are roughly evenly distributed across the four main categories.

4 POI COLLECTION

This section describes the process of creating our POI collection. Table 5 shows the information associated with an example POI in our collection.

Table 4: Number of information needs per main category.

Category	#InfoNeeds
[AL] Active Life	27
[AE] Arts & Entertainment	30
[RF] Restaurants and Food	30
[NL] Nightlife	25
Total	112

Table 5: Example POI (ID=11865).

URL	https://www.yelp.com/biz...
Image_url	https://s3-media2.fl.yelpcdn.com/bphoto/...
Name	L'Huitrier
Alias	l-huitrier-paris-2
Main_category	Restaurants and Food
Sub_categories	French, Seafood
Rating	4.0
Price	€€€
Address	16 rue Saussier Leroy
City	Paris
City_from_postal	Paris
Zip_code	75017
Country_code	FR
State_code	75
Phone	+33140548344
Review_count	22
<i>Snippets</i>	
URL	https://www.tripadvisor.com/...
Title	L'Huitrier, Paris - Ternes - Menu, Prices & Restaurant
Snippet	Order food online at L'Huitrier, Paris with Tripadvisor: See 320 unbiased reviews of L'Huitrier, ranked #1,386 on Tripadvisor among 17,679 restaurants in Paris.
URL	https://www.viamichelin.com/...
Title	L'Huitrier - Paris : a Michelin Guide restaurant
Snippet	L'Huitrier - a restaurant from the Michelin Guide Find all the information you need on the L'Huitrier restaurant: Michelin guide review, user reviews, cuisine, opening times, meal prices... L'Huitrier - a Michelin Guide restaurant.
...	

4.1 Collecting POIs from Yelp

POIs were collected from Yelp, using their Fusion API (which is permitted for non-commercial uses⁶). The POIs were primarily categorized within one of the four main categories mentioned in Section 3.1, with the exception of a few manual recommendations, see Section 5.1. The resulting collection consists of 695,053 POIs distributed across 23 countries. A breakdown of POIs per country and per main category is provided in the online repository.⁷

⁶https://www.Yelp.com/developers/api_terms

⁷<https://github.com/iai-group/sigir2021-pointrec/blob/master/Statistics.md>

4.2 Linking to Geographical Knowledge Bases

Two different geographical knowledge bases were used: Restcountries⁸ and GeoNames.⁹ Restcountries was used to retrieve the alpha2codes for every country in the world. Then, using GeoNames, we identified the cities¹⁰ within the countries present in the set of information needs.

4.3 Enrichment with Web Snippets

We used Bing's Web Search API¹¹ to further enrich a subset of the POI collection with short descriptions for each POI. Specifically, this subset is comprised of all cities that are present in our set of information needs. As a query, we combined the POI name and city, to ensure that the web snippets corresponded to the correct POI.

4.4 Data Cleaning and Post-processing

The returned data from Yelp would occasionally include a very specific location as the city, e.g., a neighborhood within a city. To link these locations to their appropriate city name, we used the Zip_code field together with GeoNames' postal code table. However, even the Zip_code field of the returned data would sometimes be inconsistent.¹²

Additionally, some POIs were missing their main category. These were linked to their corresponding main category, by matching the sub-categories between the POI and the main categories.

5 RECOMMENDATIONS AND RELEVANCE ASSESSMENTS

We follow the standard practice of pooling [20] for collecting relevance annotations. For each information need, a pool of candidate POIs is constructed by combining recommendations made by humans as well as by multiple configurations of a simple content-based recommender system. These would be termed "manual" and "automatic" runs, respectively, following TREC terminology.

5.1 Manual Recommendations

Manual recommendations are solicited from two sources: extracted from forum posts and gathered via crowdsourcing.

5.1.1 Recommendations from Forum Posts. Given that information needs were collected from community QA sites, we can tap into the respective forum threads as a source of recommendations. Specifically, two authors of the paper manually processed the replies to the originating post of each information need to collect POI suggestions.

Note that some posts had no replies, thus no suggestions. In fact, this was the case for 31 out of 112 information needs, with most of these originating from Yahoo! Answers. On the contrary, some posts had too many replies. Due to time restrictions, the five most relevant comments (following the default relevance-based ranking of the service) were used to gather POI suggestions. The next step

⁸<https://restcountries.eu>

⁹<https://www.geonames.org>

¹⁰Our definition of a city requires there to be at least 1,000 inhabitants in the given area. As a result, 138,088 cities were identified.

¹¹<https://azure.microsoft.com/nb-no/services/cognitive-services/bing-web-search-api>

¹²In most cases however, at least one of these fields were correct.

Valentine's Day Dinner Recommendations

My wife and I will be in Paris on Valentine's Day and want to have a nice romantic meal somewhere. Our budget is probably around €200. Where would you recommend? EDIT: Usually in the US, a romantic dinner for us is seafood (my wife loves scallops) or a nice steak at a more upscale restaurant. As far as type of food, I think either some french cuisine, seafood or steak.

[Link to Original](#)

Metadata

City Country Group
Paris FR My wife and I

Main Category
Restaurants and Food

Sub categories
French, Seafood, Steakhouse

Constraints
Must Romantic dinner, scallops

Should French cuisine, seafood, steak, upscale restaurant, €200 budget

Find a good match to the request

Look at the request, and find a Yelp web-address to a Point of Interest (POI) that satisfies the requirements. It is important that you consider the location (city, country and state), category and the constraints that are highlighted below the request.

Note that the link you provide, has to be a yelp business link, i.e., it has to be in the form [https://yelp.com/biz/...](https://yelp.com/biz/)

You need to provide at least one Yelp POI. For every additional POI you find (up to 4), you'll get a \$ 0.10 bonus.

If you cannot find any POI for the given location, then you may provide the link to the POI at the closest location you can find in the text area, you'll need to explain how you arrived at this result.

Links to the points if interest
[https://yelp.com/biz/...](https://yelp.com/biz/)

[https://yelp.com/biz/...](https://yelp.com/biz/)

[https://yelp.com/biz/...](https://yelp.com/biz/)

[https://yelp.com/biz/...](https://yelp.com/biz/)

[https://yelp.com/biz/...](https://yelp.com/biz/)

Describe how you found the point if interest
I searched...

Figure 1: Crowdsourcing user interface for collecting POI recommendations.

was to find the corresponding Yelp page, to collect the Yelp POI link, which was used later in a similar fashion to the links collected through crowdsourcing, see Section 5.1.2.

5.1.2 Recommendations Collected via Crowdsourcing. Crowd workers, using the Amazon Mechanical Turk platform, were tasked with providing POI suggestions for a given information need. Specifically, a user interface was developed (and refined in multiple iterations via pilot runs), where the most important metadata fields associated with the information need are highlighted; see Figure 1.

Each worker was tasked to find at least 1 and at most 5 POIs. The base payment was \$0.4, with a bonus of \$0.1 for each additional POI provided (on top of the mandatory one). Since we aimed to collect as many suggestions as possible, each information need was processed by 5 different workers. Workers were required to have more than 1,000 approved HITs, and a HIT approval rate above 95%. The submissions were manually checked by the paper authors.

5.2 Automatic Recommendations

Additionally, a simple content-based recommender system was implemented. Specifically, the POI collection was indexed using Elasticsearch, using a separate field for each POI attribute. Then, three baseline rankers, incorporating various filtering criteria, were implemented by using structured queries. The main differences between these methods are summarized in Table 6; all three give a reward to POIs with a matching city.

- **Baseline 1:** The main category of the POI was required to match the information need; categories have also been exploited in prior work to reduce the pool of candidates, e.g., [15, 19, 28]. Additionally, POIs that match sub-categories will receive a higher relevance score.
- **Baseline 2:** This method does not take into account main and sub-categories. Instead, it uses the request and constraints specified by the user and matches those against the text collected from web snippets.
- **Baseline 3:** It is a combination of Baselines 1 and 2, and thus represents the strictest approach.

Table 6: Automatic recommender systems.

System	Categories	Const. & Reqs.
Baseline 1	+	
Baseline 2		+
Baseline 3	+	+

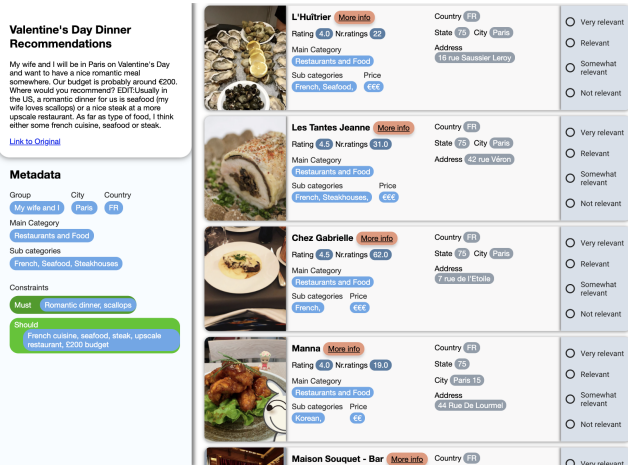


Figure 2: Crowdsourcing UI for relevance assessments.

5.3 Relevance Labeling

Relevance assessments were collected using crowdsourcing, using the interface shown in Figure 2. Each task consisted of 5 POI recommendations for a given information need and paid \$0.15. POI recommendations were pooled together from manual suggestions and the three automatic baselines (with up to top-50 results from each baseline). The order of POIs was randomized to eliminate potential position bias.

Each information need-POI pair was judged by at least 3 workers on a 4-point relevance scale:

- A POI is *highly relevant* (3), if it perfectly matches the request, with respect to location, category, and constraints.
- A POI is *relevant* (2), if the location and category match, and the requester would likely be interested in it.
- A POI is *somewhat relevant* (1), if the location matches, and the requester could possibly be interested in it.
- Otherwise, the POI is *not relevant* (0).

In case of strong disagreement, that is, 2 points difference on the relevance scale, 2 additional judgments were collected. Relevance labels were then consolidated by taking a majority vote.

Worker qualifications were the same as in Section 5.1.2. To ensure data quality and to avoid spammers, we used both gold answers (POI suggestions extracted manually from forums by the paper authors) and honeypots (injecting erroneous suggestions from a different main category). Workers with an error rate above 30% were rejected. The mean task completion time was 261sec.

6 BASELINE RESULTS

This section presents an experimental comparison of the baselines. Our main measure is NDCG@5 (following the TREC Contextual

Table 7: Baseline results. Highest scores are boldfaced.

System	NDCG@5	NDCG@10	MRR	MAP
Baseline 1	0.6389	0.5812	0.5812	0.3304
Baseline 2	0.4109	0.3979	0.2814	0.0667
Baseline 3	0.6784	0.6573	0.5535	0.2506




Figure 3: NDCG@5 difference between Baseline 1 and Baseline 3. Each bar corresponds to an information need. Positive values indicate that Baseline 1 performs better, negative values indicate that Baseline 3 is more effective.

Suggestions track [9, 12]); we also report on NDCG@10. Additionally, we consider binary evaluation metrics, MAP and MRR, accepting only the highest relevance level (3) as relevant.

Table 7 displays the results. We observe that the strictest method (Baseline 3) performs best in terms of graded relevance measures, while Baseline 1 is the most effective in terms of binary metrics. Figure 3 shows a comparison between Baselines 1 and 3 on the level of individual information needs. We find that on about half of the information needs Baseline 1 works better, while on the other half Baseline 3 is preferred. Recall that both methods consider categories, while Baseline 3 additionally matches the requests and constraints as well against web snippets. These results show that constraints can potentially aid retrieval, but more sophisticated approaches are needed in order to effectively incorporate them.

7 CONCLUSION AND FUTURE WORK

This paper has introduced the POINTREC test collection for contextual POI recommendation. It consists of a collection of POIs across 23 countries, a set of 112 information needs originating from real users, and graded relevance assessments collected via crowdsourcing. Additionally, we have developed and experimentally compared three variants of a simple retrieval-based baseline approach.

This collection is primarily aimed at facilitating research on POI recommendation in a narrative-based scenario, where preferences and constraints are given in natural language. These have been manually identified and provided as metadata annotations, and have the potential to improve the quality of recommendations. The automatic extraction, representation, and interpretation of user preferences is an interesting and challenging natural language understanding task on its own account that is not exclusive to our setting; it is also a central component, for instance, in the context of conversational assistants (see, e.g., [2]). Our annotated set of information needs provides a resource for research in this area.

REFERENCES

- [1] Mohammad Aliannejadi and Fabio Crestani. 2018. Personalized Context-Aware Point of Interest Recommendation. *ACM Trans. Inf. Syst.* 36, 4, Article 45 (Oct. 2018).
- [2] Krisztian Balog, Filip Radlinski, and Alexandros Karatzoglou. 2021. On Interpretation and Measurement of Soft Attributes for Recommendation. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '21)*.
- [3] Jan R. Benetka, Krisztian Balog, and Kjetil Nørveg. 2017. Anticipating Information Needs Based on Check-in Activity. In *Proceedings of the 10th ACM International Conference on Web Search and Data Mining (WSDM '17)*. 41–50.
- [4] Toine Bogers and Marijn Koolen. 2017. Defining and Supporting Narrative-Driven Recommendation. In *Proceedings of the 11th ACM Conference on Recommender Systems (RecSys '17)*. 238–242.
- [5] Matthias Braunhofer, Francesco Ricci, Béatrice Lamche, and Wolfgang Wörndl. 2015. A Context-Aware Model for Proactive Recommender Systems in the Tourism Domain. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct (MobileHCI '15)*. 1070–1075.
- [6] Chen Cheng, Haiqin Yang, Irwin King, and Michael R. Lyu. 2012. Fused Matrix Factorization with Geographical and Social Influence in Location-Based Social Networks. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI '12)*. 17–23.
- [7] Chen Cheng, Haiqin Yang, Michael R. Lyu, and Irwin King. 2013. Where You like to Go next: Successive Point-of-Interest Recommendation. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI '13)*. 2605–2611.
- [8] Eunjoon Cho, Seth A. Myers, and Jure Leskovec. 2011. Friendship and Mobility: User Movement in Location-Based Social Networks. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '11)*. 1082–1090.
- [9] Adriel Dean-Hall, Charles L. A. Clarke, Jaap Kamps, Julia Kiseleva, and Ellen M. Voorhees. 2015. Overview of the TREC 2015 Contextual Suggestion Track. In *Proceedings of the 24th Text Retrieval Conference (TREC '15)*.
- [10] Huiji Gao, Jiliang Tang, Xia Hu, and Huan Liu. 2013. Exploring Temporal Effects for Location Recommendation on Location-Based Social Networks. In *Proceedings of the 7th ACM Conference on Recommender Systems (RecSys '13)*. 93–100.
- [11] Huiji Gao, Jiliang Tang, Xia Hu, and Huan Liu. 2015. Content-Aware Point of Interest Recommendation on Location-Based Social Networks. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI '15)*. 1721–1727.
- [12] Seyyed Hadi Hashemi, Jaap Kamps, Julia Kiseleva, Charles L. A. Clarke, and Ellen M. Voorhees. 2016. Overview of the TREC 2016 Contextual Suggestion Track. In *Proceedings of the 25th Text Retrieval Conference (TREC '16)*.
- [13] Defu Lian, Cong Zhao, Xing Xie, Guangzhong Sun, Enhong Chen, and Yong Rui. 2014. GeoMF: Joint Geographical Modeling and Matrix Factorization for Point-of-Interest Recommendation. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '14)*. 831–840.
- [14] Defu Lian, Kai Zheng, Yong Ge, Longbing Cao, Enhong Chen, and Xing Xie. 2018. GeoMF++: Scalable Location Recommendation via Joint Geographical Modeling and Matrix Factorization. *ACM Trans. Inf. Syst.* 36, 3, Article 33 (March 2018).
- [15] Xin Liu, Yong Liu, Karl Aberer, and Chunyan Miao. 2013. Personalized Point-of-Interest Recommendation by Mining Users' Preference Transition. In *Proceedings of the 22nd ACM International Conference on Information & Knowledge Management (CIKM '13)*. 733–738.
- [16] Yiding Liu, Tuan-Anh Nguyen Pham, Gao Cong, and Quan Yuan. 2017. An Experimental Evaluation of Point-of-Interest Recommendation in Location-Based Social Networks. *Proc. VLDB Endow.* 10, 10 (June 2017), 1010–1021.
- [17] Hossein A. Rahmani, Mohammad Aliannejadi, Mitra Baratchi, and Fabio Crestani. 2020. Joint Geographical and Temporal Modeling Based on Matrix Factorization for Point-of-Interest Recommendation. In *Proceedings of the 42nd European Conference on Information Retrieval (ECIR '20)*. 205–219.
- [18] Hossein A. Rahmani, Mohammad Aliannejadi, Rasoul Mirzaei Zadeh, Mitra Baratchi, Mohsen Afsharchi, and Fabio Crestani. 2019. Category-Aware Location Embedding for Point-of-Interest Recommendation. In *Proceedings of the 2019 ACM SIGIR International Conference on Theory of Information Retrieval (ICTIR '19)*. 173–176.
- [19] Jitao Sang, Tao Mei, and Changsheng Xu. 2015. Activity Sensor: Check-In Usage Mining for Local Recommendation. *ACM Trans. Intell. Syst. Technol.* 6, 3, Article 41 (April 2015).
- [20] Ellen M. Voorhees and Donna K. Harman. 2005. *TREC: Experiment and Evaluation in Information Retrieval (Digital Libraries and Electronic Publishing)*. The MIT Press.
- [21] Hao Wang, Huawei Shen, Wentao Ouyang, and Xueqi Cheng. 2018. Exploiting POI-Specific Geographical Influence for Point-of-Interest Recommendation. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI '18)*. 3877–3883.
- [22] Dingqi Yang, Bingqing Qu, Jie Yang, and Philippe Cudre-Mauroux. 2019. Revisiting User Mobility and Social Relationships in LBSNs: A Hypergraph Embedding Approach. In *The World Wide Web Conference (WWW '19)*. 2147–2157.
- [23] Dingqi Yang, Bingqing Qu, Jie Yang, and Philippe Cudre-Mauroux. 2020. LBSN2Vec++: Heterogeneous Hypergraph Embedding for Location-Based Social Networks. *IEEE Transactions on Knowledge and Data Engineering* (2020).
- [24] Dingqi Yang, Daqing Zhang, and Bingqing Qu. 2016. Participatory Cultural Mapping Based on Collective Behavior Data in Location-Based Social Networks. *ACM Trans. Intell. Syst. Technol.* 7, 3, Article 30 (Jan. 2016).
- [25] Mao Ye, Peifeng Yin, Wang-Chien Lee, and Dik-Lun Lee. 2011. Exploiting Geographical Influence for Collaborative Point-of-Interest Recommendation. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '11)*. 325–334.
- [26] Hongzhi Yin, Bin Cui, Xiaofang Zhou, Weiqing Wang, Zi Huang, and Shazia Sadiq. 2016. Joint Modeling of User Check-in Behaviors for Real-Time Point-of-Interest Recommendation. *ACM Trans. Inf. Syst.* 35, 2, Article 11 (Oct. 2016).
- [27] Quan Yuan, Gao Cong, Zongyang Ma, Aixun Sun, and Nadia Magnenat Thalmann. 2013. Time-Aware Point-of-Interest Recommendation. In *Proceedings of the 36th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '13)*. 363–372.
- [28] Jia-Dong Zhang and Chi-Yin Chow. 2015. GeoSoCa: Exploiting Geographical, Social and Categorical Correlations for Point-of-Interest Recommendations. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '15)*. 443–452.