Impact of Semantic Granularity on Geographic Information Search Support

(Article begins on next page)

19 April 2024

Impact of Semantic Granularity on Geographic Information Search Support

N. Mauro, L. Ardissono
Computer Science Department
University of Torino
Torino, Italy
noemi.mauro@unito.it,
liliana.ardissono@unito.it

L. Di Rocco, G. Guerrini
DIBRIS
University of Genova
Genova, Italy
laura.dirocco@dibris.unige.it,
giovanna.guerrini@unige.it

M. Bertolotto
School of Computer Science
University College Dublin
Dublin, Ireland
michela.bertolotto@ucd.ie

Abstract—The Information Retrieval research has used semantics to provide accurate search results, but the analysis of conceptual abstraction has mainly focused on information integration. We consider session-based query expansion in Geographical Information Retrieval, and investigate the impact of semantic granularity (i.e., specificity of concepts representation) on the suggestion of relevant types of information to search for. We study how different levels of detail in knowledge representation influence the capability of guiding the user in the exploration of a complex information space. A comparative analysis of the performance of a query expansion model, using three spatial ontologies defined at different semantic granularity levels, reveals that a fine-grained representation enhances recall. However, precision depends on how closely the ontologies match the way people conceptualize and verbally describe the geographic space.

Index Terms—geographical information retrieval, semantic granularity, session-based concept suggestion

I. INTRODUCTION

Several researchers have used semantic knowledge to support query expansion and reformulation in information search support; e.g., [1]. Moreover, in Ontology-Driven Geographic Information Systems, abstraction has been analyzed to understand its "potential for information retrieval at different levels of granularity" [2]. However, to the best of our knowledge, the impact of different domain conceptualizations on the accuracy of information search support has not been fully investigated yet. We focus on Geographical Information Retrieval, and aim at studying the influence of semantic granularity (i.e., the degree of specificity in concept representation) on query expansion. Specifically, we aim at measuring the extent to which, by combining general information about search behavior with a more generic, or a more detailed domain conceptualization,

an automated system can learn regularities useful in the suggestion of concepts relevant to the user's information needs. For instance, by analyzing general search behavior we might discover that, if a user looks for kindergartens in a town, he or she might also be interested in information related to other children activities, such as play or sport areas.

Notice that our aim is not only to measure the precision of the system's suggestions, but also to define a notion of "richness" based on the number of relevant suggested concepts, as this is important for catalog exploration. Specifically, we investigate the following research questions:

RQ1: *What is the relationship between the semantic granularity of a domain conceptualization and the accuracy in suggesting types of information relevant to the user's needs during an exploratory search task?*

RQ2: *Does semantic granularity in the representation of geographical knowledge influence the richness of concept suggestion?*

In order to answer these questions, we need to know which concepts people frequently focus on during a search session, as a source of evidence of co-occurring information needs. Thus, we apply a session-based concept suggestion model to extract concepts co-occurrence, and we compare its performance by using different ontologies to train it and to interpret the search queries: a fine-grained ontology, a less detailed one developed for OpenStreetMap (www.openstreetmap.org) and the GeoNames Mappings ontology (www.geonames.org/ontology/mappings_v3.01.rdf). The results of an experiment based on a large query log reveal that a finer-grained semantic granularity improves recall and richness of results: with the first ontology the concept suggestion model achieves the best recall and supports the generation of more suggestions than with the other two. However, the model achieves the best precision and accuracy (F1) using the second ontology, which is based on crowdsourced data.

In the following, Section II summarizes the related research. Sections III and IV present the ontologies and the concept suggestion model we employed. Section V describes the datasets we used. Section VI describes our empirical evaluation. Section VII concludes the paper.
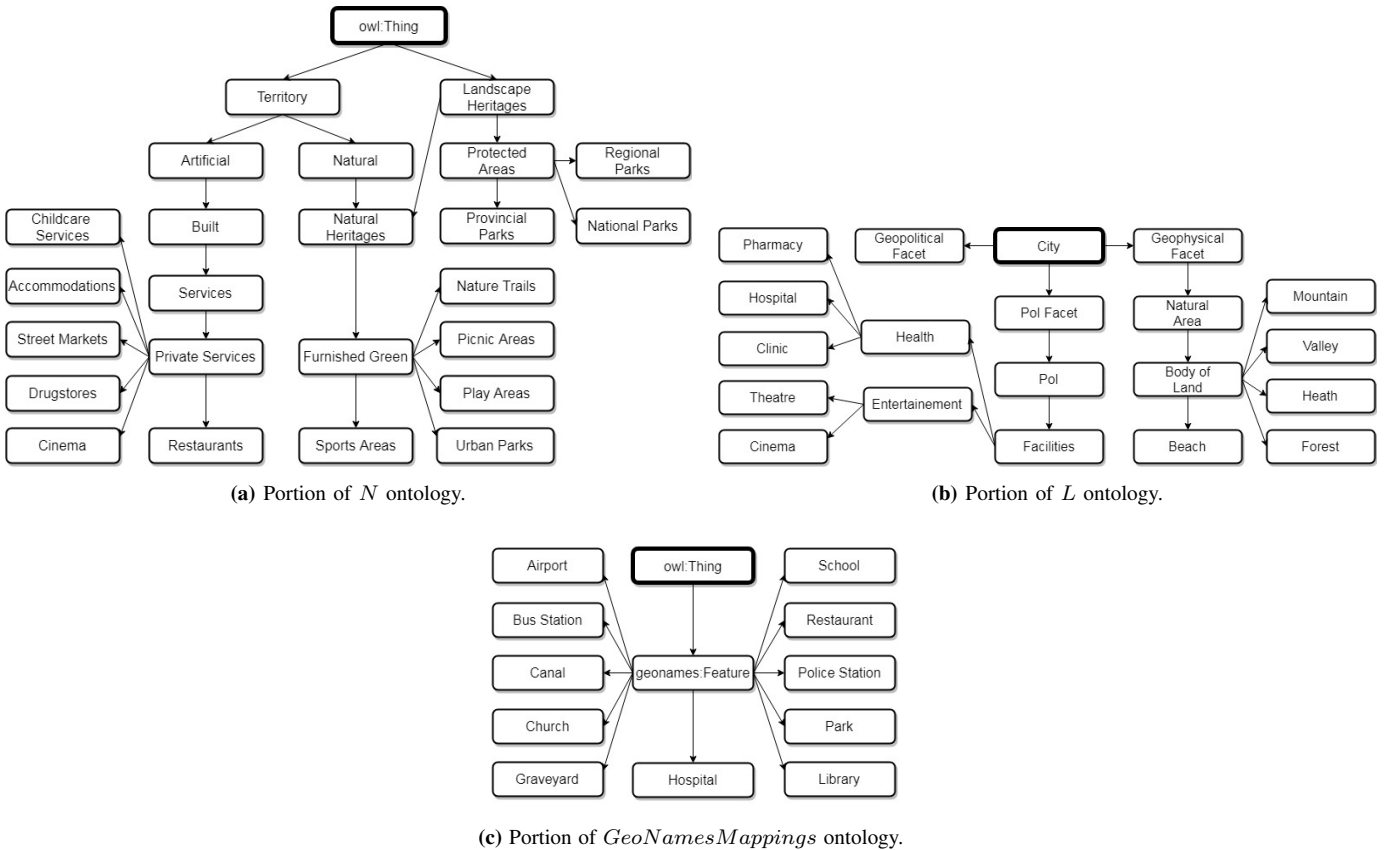
(a) Portion of N ontology.

(b) Portion of L ontology.

(c) Portion of $GeoNamesMappings$ ontology.

Fig. 1: Ontologies selected for the empirical evaluation. The root nodes of the ontologies have thick borders for readability purposes.

II. RELATED WORK

Geographical Information Retrieval mainly focuses on the interpretation of the spatial component of a search task [3], [4]. However, Ballatore et al. [5] identify research themes and open questions, which include aspects related to search behavior models and semantic aspects of spatial search. In this context, we focus on geospatial knowledge, and aim to identify the potential of semantic granularity in search support.

Term suggestion and query expansion have been based on the analysis of the lexical or semantic distance among terms [6], [7], on term co-occurrence in search logs [8], [9] and documents [10], and on an analysis of click distribution on search results [11]. Jiang et al. propose a framework for search session derivation, geographic information extraction and geographic web search topic discovery, which exploits topic modeling techniques [12]. Differently, we aim at proposing concepts by using common patterns of concept exploration.

In Geographic Information Retrieval, ontologies have been used to improve geographic *features* extraction [13]. On the other hand, we employ them for geographic *concepts* extraction, in order to provide the user with topics to explore, rather than individual items.

Researchers developed specific geographic ontologies, such as GeoNames. Moreover, they attempted to semantify geographical data targeted to specific tasks; e.g., the LinkedGeoData ontology links OpenStreetMap information to DBpedia,

GeoNames, and others ontologies [14]. Fonseca et al. propose an ontology to classify geographic elements with respect to geometrical features and attribute values; i.e., semantic features [15]. Finally, a relevant amount of research has been devoted to semantifying geospatial Open Data and crowdsourced data. For instance, OSMonto [16] attempts to structure the shallow implicit ontology of OpenStreetMaps tags. Moreover, in [17] a Semantic Web resource extracted from the OpenStreetMap Wiki website is proposed, to semantify crowdsourced data.

III. THE USED ONTOLOGIES

In this section we analyze and compare the three geographical ontologies used for the experiments. We focus on:

1) Semantic granularity, i.e., the level of detail at which geographic objects are described [2]. This is different from spatial granularity, which refers to the granularity of geographic toponyms; e.g., see [3].
2) Ontology structure, quantified in terms of number of concepts and number of subclass relations.

N ontology (see Figure 1a) contains fine-grained concepts about cities [18]. This ontology provides a multi-faceted specification of classes and relations that characterize the information space considering three high-level aspects: natural (e.g., parks), artificial (e.g., infrastructures) and normative (e.g., administrative boundaries). These aspects are specialized into classes at more than one level of detail. For instance, as
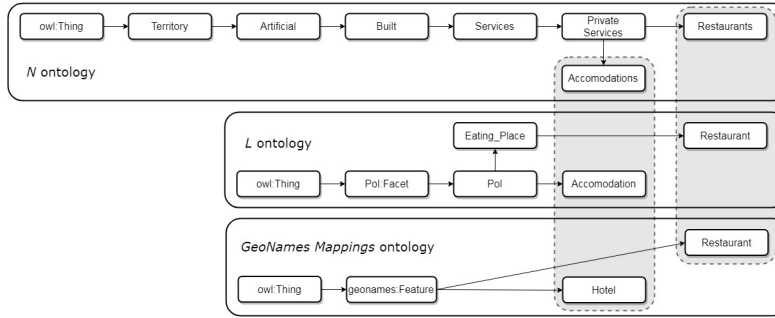
Fig. 2: Portion of the three ontologies with an example of shared concepts (in gray areas).

TABLE I: Structural comparison of the GeoNames, L and N ontologies.

Metric	N	L	GeoNames
Total number of classes	195	97	150
Total number of subclass relations	268	94	155
Longest path to leaves (graph nodes)	6	5	3
Number of used classes	72	70	136
Number of used subclass relations	81	69	135
Longest path to used leaves (graph nodes)	6	5	2
Mean degree of used classes (graph nodes)	4.5	3.63	67.5

partially reported in the figure, Landscape Heritage is specified into a deep hierarchy of classes in order to distinguish various types of Furnished Green from Protected Areas. The longest path between owl:Thing and the leaf classes has length = 6.

L ontology (see Figure 1b) structures OpenStreetMap tags that refer to geographic concepts in the context of cities [19]. It is multi-faceted and has a similar structure, but coarser granularity, than N; moreover, being derived from the largely-used OpenStreetMap repository, it reflects a conceptualization that comes from the general public. The geographic information about a city is described using three facets: Point of Interest (PoI), geoPolitical and geoPhysical. The longest path between the owl:Thing class and the leaf classes of L has length = 5.

GeoNames Mappings ontology (see Figure 1c) maps the GeoNames taxonomy to external ontologies, such as schema.org and linkedgeodata.org. GeoNames derives from the geonames.org database and describes more than 9,000,000 features categorized in 645 alphanumeric feature codes. The Mappings ontology maps 108 of its codes and describes knowledge at a coarse semantic granularity: information is organized in a flat structure including a top-level OWL class (owl:Thing), and a child class (geonames:Feature) that has all the GeoNames codes as its children. Only two subclasses of geonames:Feature, representing toponyms, have children themselves. Thus, even though the longest path between owl:Thing and the leaves of the ontology has length = 3, the path from owl:Thing to most classes has length = 2.

As a reference for the comparison among ontologies, we selected GeoNames Mappings (henceforth, GeoNames for brevity), which is largely adopted in the geospatial community. In each ontology, we only considered geospatial classes related to natural features of the territory (e.g, rivers, mountains) and artificial features (e.g., malls, buildings, streets). We excluded

administrative boundaries, which concern spatial, rather than semantic, granularity and are abstract concepts. The ontologies used in our analysis are thus subsets of the original ones.

Table I shows a set of metrics to describe the complete ontologies, and the subsets focused on natural and artificial features. As shown in row "Number of used classes", the used subset of GeoNames is the largest one, followed by the N and the L ones. The three ontologies share 8 concepts; GeoNames and L share 23 ones; GeoNames and N share 17 ones.

The subsets of N, L and GeoNames have different semantic granularity. A rough measure of this can be obtained by jointly considering the length of the longest path between owl:Thing and the leaves of an ontology, and the mean degree of the nodes of the ontology graph. Among the three, N has the maximum longest path (length = 6) and its mean degree is 4.5, which denotes the fact that several classes are specified into a relatively small set of subclasses. L has few subclasses per node as well (mean degree = 3.63) but it defines concepts at a slightly more superficial level (length of longest path = 5). Finally, GeoNames is flat (length = 2) and has a very high mean node degree (67.5). We report some examples to compare the ontologies. Class Park in GeoNames is represented in N as Regional_Park, Provincial_Park and National_Park. Therefore, we can say that Park is shared between N and GeoNames but they differ in granularity and thus the two representations can support concept identification at different levels of detail. Similarly, Shopping of L corresponds to Shopping_mall, Market and Shop in N. Finally, in N, concept Restaurant is associated with a much longer and descriptive path than in GeoNames. The corresponding path in L is descriptive but not as exhaustive as in N; see Figure 2.

IV. CONCEPT SUGGESTION MODEL (CS MODEL)

In order to compare suggestion accuracy using different reference ontologies, we used the Concept Suggestion Model described in [20]. This model suits our needs because it extracts concept co-occurrence (w.r.t. term, or query co-occurrence, returned by other algorithms) from the search sessions of a query log, and thus enables us to identify information needs from a semantic point of view. The CS model bases concept suggestion on (i) the definition of co-occurrence clusters

grouping concepts that are frequently referenced together in search sessions; (ii) the adoption of strategies that suggest concepts by matching the co-occurrence clusters to the search context according to different policies. By search context we mean the set of concepts referenced in the observed queries.

A. Creation of Concept Co-occurrence Clusters

The process to create the concept co-occurrence clusters, starting from the query log and the ontology conceptualizing the domain knowledge, includes the following steps:

1) Split the query log in sessions by aggregating queries according to their temporal proximity.
2) Produce a *reduced dataset* by selecting the search sessions whose queries reference at least one concept corresponding to an ontology class. For this task, the ontology classes are annotated with linguistic and domain knowledge. The lemmas of the words occurring in the queries are thus matched to the lemmas of the class annotations in order to find correspondences.
3) Build a concepts co-occurrence graph by analyzing the sessions of the reduced dataset. In the graph, nodes denote concepts, edges denote concept co-occurrence in search sessions, and the weight of an edge denotes co-occurrence frequency of its input and output concepts. In order to focus on common search behavior, the edges having low weight are pruned.
4) Apply a community detection algorithm (COPRA [21]) to the graph to identify its concept communities.

The output of this phase is a set of clusters, each one consisting of an unordered set $CL = \{c_1, \ldots, c_n\}$ of concepts c_x that highly co-occur in the sessions of the reduced dataset.

B. Concept Suggestion Strategies

Given a search session $S = \{Q_1, \ldots, Q_i, \ldots, Q_n\}$, let $S@i$ be the observed portion of S, i.e., $\{Q_1, \ldots, Q_i\}$. Moreover, let $C@i = \{c_1, \ldots, c_k\}$ be the set of concepts identified by interpreting the queries of $S@i$.

In [20], three main strategies for cluster selection are proposed, which result in different concept suggestion policies:

SLACK: select a set of clusters $\{CL_{x_1}, \ldots, CL_{x_y}\}$ that contain *at least* one concept c such that $c \in C@i$; i.e., c has been referenced in the observed portion of the session ($S@i$). For instance, given $C@i = \{c_1, c_3\}$, and three clusters, $CL_1 = \{c_1, c_3, c_7\}$, $CL_2 = \{c_2, c_3, c_5, c_8\}$, $CL_3 = \{c_5, c_8, c_9\}$, SLACK selects CL_1 and CL_2, and it suggests $\{c_2, c_5, c_7, c_8\}$.

SLACK SELECTIVE: same as SLACK, but only one cluster that best matches $C@i$ is selected. We compute the degree of matching between a cluster CL and $C@i$ as $|CL \cap C@i|$. In the example, the strategy suggests $\{c_7\}$ because the best matching cluster is CL_1.

STRICT: select the clusters $\{CL_{x_1}, \ldots, CL_{x_y}\}$ containing *all* the concepts $c \in C@i$. In the example, the strategy does not suggest any concepts.

We aim at generating suggestions immediately after the first query of a search session; thus, we apply the strategies to $S@1$.

TABLE II: Size of *AOL-reduced* datasets.

Measure	AOL_N	AOL_L	AOL_{GN}
Number of queries	1,581,817	1,486,122	1,443,448
Number of sessions	945,945	911,399	864,869
Number of users	247,868	240,474	232,931

V. DATASETS FOR TRAINING THE CS MODEL

We used the AOL query log[1] as a source of search sessions. Although that log was involved in an information leak issue, we decided to use it for two reasons: firstly, our analysis is ethically correct, because we model general search behavior. The preference co-occurrence clusters we extract represent aggregate data about people behavior, and they abstract from the search histories of particular users. Secondly, to the best of our knowledge, the AOL log is the only public, large dataset that reports detailed data about textual search queries, and that can be used for linguistic interpretation.[2]

Each line of the AOL log represents a submitted query, or a click-through event, and reports an anonymous user id, the search query, date and hour when the query was submitted and other fields that are not relevant here.

A. AOL-reduced Datasets

To compare the impact of domain conceptualization on concept suggestion, we reduced the AOL query log using the N, L and GeoNames ontologies. We generated three *AOL-reduced* datasets (AOL_N, AOL_L, AOL_{GN}) by filtering the sessions of the AOL log according to the concepts defined in each ontology; see Section IV-A. For this task, we pre-processed the ontologies by automatically annotating their classes with linguistic knowledge (synonyms and linguistic definitions), using BabelFy multilingual Entity Linking and Word Sense Disambiguation service [22], and Stanford CoreNLP lemmatizer [23].

As shown in Table II, the AOL_N dataset contains the highest number of queries, sessions and users. AOL_L contains the second highest number of queries, sessions and users, and AOL_{GN} is the smallest dataset, even though the GeoNames ontology is larger than the other two. We explain this finding using semantic granularity: N is the most detailed ontology and supports the selection of the largest number of queries from the log. However, in GeoNames, the coarse-grained classes limit the identification of concepts in search queries, and make the selection of relevant sessions more constrained than using L. The datasets have some similarities. As reported in Table III, the mean length of their sessions is almost the same: 1.67 queries in AOL_N and AOL_{GN}, 1.63 in AOL_L. Moreover, the distribution of sessions w.r.t. their length follows a Power Law (most sessions are short, few contain many

---

[1]Retrieved from https://archive.org/details/AOL_search_data_leak_2006.
[2]We analyzed some public datasets but they did not meet our requirements. E.g., the Excite query dataset (https://svn.apache.org/repos/asf/pig/trunk/tutorial/data/) contains fewer queries (about 1M queries against the 20M of AOL log). In the Yahoo dataset (https://webscope.sandbox.yahoo.com/catalog.php?datatype=l&did=50) queries are coded; thus, it is not possible to extract any linguistic information to learn the concept co-occurrence clusters.

TABLE III: Length of the search sessions of the *AOL-reduced* datasets.

Measure	AOL_N	AOL_L	AOL_{GN}
Min number of queries per session	1	1	1
Max number of queries per session	93	96	103
Mean number of queries per session	1.67	1.63	1.67
Median number of queries per session	1	1	1
Standard deviation	1.45	1.38	1.44

TABLE IV: Sample concept co-occurrence clusters.

Dataset	Sample clusters
AOL_N	{Play Areas, Sport Areas, Kindergartens, Law Enforcement, Hospitals} {Play Areas, Sport Areas, Libraries, Childcare services, Law Enforcement, Schools}
AOL_L	{Educational_institution, Kindergarten, University} {Bar, Cafe, Eating_place, Pub, Restaurant}
AOL_{GN}	{Library, Museum, Police_Station, University} {Beach, Resort}

queries). Thus, we can use them to compare the concept suggestion strategies.

B. Concept Co-occurrence Clusters

For each ontology X, and AOL-reduced dataset AOL_X, we created a concept co-occurrence graph G_X as described in Section IV-A. Also in this case, we pruned G_X to remove the weakest co-occurrence relations among concepts. Then, we applied the COPRA algorithm to G_X to identify the communities of concepts that frequently co-occur in the search sessions. The algorithm returned 23 clusters for AOL_N, 19 for AOL_L and 23 for AOL_{GN}. Table IV shows two sample clusters for each dataset.

The clusters generated from AOL_N contain a larger number of concepts than the others (about 21% more than those produced using AOL_L, and about 31% more than those derived from AOL_{GN}). This finding is in line with the granularity of the ontologies: being more specific, N makes it possible to identify a more diverse set of co-occurring concepts in the search sessions. In contrast, GeoNames supports the identification of few concepts, and thus produces small clusters.

The clusters generated using AOL_L are more topic-centered than the other ones: the concepts of each AOL_L cluster have the same parent concept in the L ontology. As L is fairly fine-grained, it describes different topics as non-leaf concepts of the ontology. Thus, having the same parent means being strictly related from a semantic point of view. Indeed, AOL_N generates a few topic-centered clusters, as well. However, most of its clusters include semantically distant concepts that typically descend from different high-level concepts of the ontology. For instance, in Table IV, Play Areas and Sport Areas descend from Natural, a high-level class of the N ontology. The other concepts of the cluster descend from Artificial (another high-level class). Finally, the clusters generated from AOL_{GN} cover a broad range of topics, given the flat structure of the GeoNames ontology.

VI. EMPIRICAL EVALUATION

We measured the performance of the CS Model on the AOL-reduced datasets, using the respective concept co-occurrence clusters for concept suggestion. For each dataset, we measured the richness and the accuracy of the suggestions that the SLACK, SLACK-SELECTIVE and STRICT strategies generated, by matching the concepts they proposed, given the first query of a session S ($S@1$), against those referenced in the remainder of S; i.e., those actually explored by users. For the evaluation we applied 10-fold cross-validation on each dataset, after having randomly distributed its sessions on folders. We used 90% of the sessions as learning set and 10% as test set.

Table V presents the evaluation results obtained by applying SLACK and SLACK-SELECTIVE to the *AOL-reduced* datasets, considering only the sessions that include at least two queries. We omit the results of STRICT due to space limitations.

Henceforth, C_S denotes the set of concepts explored by the user in a session S; $C_{Sugg_{S@i}}$ is the set of concepts suggested by a strategy after analyzing the $S@i$ portion ($i = 1$).

Richness of suggestions. The first 3 lines of Tables Va and Vb compare the numbers of relevant concepts suggested by the strategies, per dataset. These numbers measure the suggested concepts that are explored by the users, and thus represent useful suggestions. It can be seen that SLACK recommends more relevant concepts than SLACK-SELECTIVE. However, both strategies obtain the highest values when they are applied to AOL_N, which is the most detailed ontology.

Recall ($R@i$). This measure represents the percentage of concepts explored in search sessions that are suggested by the strategies. It represents the system capability to cover the topics that users actually explore in their searches. Tables Va and Vb show that strategies achieve the best $R@1$ with AOL_N.

Precision ($P@i$) represents the percentage of suggested concepts that are explored in the remainder of the sessions. Tables Va and Vb show that the strategies achieve the best $P@1$ with AOL_L, and this value is much higher than those obtained on the other datasets.

Accuracy ($F1@i$) integrates precision and recall. Both strategies obtain the highest $F1@1$ score on AOL_L. Figure 3 shows the variation of $F1@1$ w.r.t. the length of search sessions (length $>= 2$). The points denote the mean value of F1 for each length value. The black lines represent the local regression on this data and show the trend of F1: at first, it decreases; then the regression line stabilizes. For longer sessions, the values of F1 do not correspond to any trend. Thus, the confidence of the regression line (in gray) decreases.

The results of the evaluation contribute to answer our research questions. Regarding **RQ1**, results support the hypothesis of a dependency between the semantic granularity of a domain conceptualization and the accuracy in the suggestion of concepts relevant to a user search. Specifically, a finer-grained semantic granularity enhances recall by supporting the generation of larger concept co-occurrence clusters that have a broad scope (they can refer to semantically distant concepts of the ontology). Thus, given a set of concepts identified in the observed query, the selection of one or more clusters that match it provides a larger pool of concepts to be suggested.

(a) SLACK performance (S@1).

	AOL_N	AOL_L	AOL_{GN}
Min number of suggested concepts	0	0	0
Max number of suggested concepts	**7**	4	4
Mean number of suggested concepts	**2.688**	0.496	1.371
Recall	**0.791**	0.783	0.768
Precision	0.614	**0.882**	0.712
F1	0.691	**0.829**	0.739

(b) SLACK-SELECTIVE performance (S@1).

	AOL_N	AOL_L	AOL_{GN}
Min number of suggested concepts	0	0	0
Max number of suggested concepts	**5**	4	3
Mean number of suggested concepts	**2.309**	0.455	1.054
Recall	**0.783**	0.759	0.761
Precision	0.621	**0.875**	0.725
F1	0.692	**0.813**	0.743

TABLE V: Performance of the concept suggestion strategies applied to the first query (S@1) of the sessions of the *AOL-reduced* datasets.

(a) SLACK

(b) SLACK-SELECTIVE
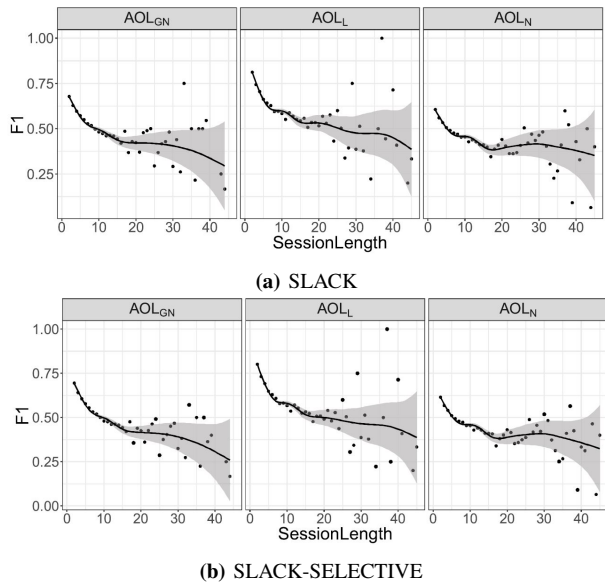
Fig. 3: Distribution of F1 score with respect to the length of sessions.

Precision increases when using an ontology that covers a broad range of concepts, like GeoNames. However, it depends on how close the ontology is to the way people conceptualize and refer to geographic information. The evaluation shows that the suggestions obtained with the crowdsourced L ontology, which is more specific than GeoNames, but less than N, are more precise than those obtained with the other ontologies.

Regarding **RQ2**, the results support the hypothesis that the semantic granularity of a domain conceptualization impacts on the richness of concept suggestions. A finer-grained ontology supports the suggestion of a larger set of relevant concepts because it is able to match a larger variety of concepts in the query logs. Moreover, it is associated with the largest clusters; thus, given a search context, it has higher suggestion capability.

VII. Conclusions

We investigated how semantic granularity in geographical knowledge conceptualization influences concept suggestions in information search support. We found that a finer-grained domain conceptualization supports the suggestion of a larger set of concepts, with higher recall. However, the relationship between the formal conceptualization (ontology) and the way people conceptualize and verbally describe geographic space influences precision and, consequently, overall accuracy.

References

[1] Y. Wang, H. Huang, and C. Feng, "Query expansion based on a feedback concept model for microblog retrieval," in *Proc. of WWW '17*, Perth, Australia, 2017, pp. 559–568.

[2] F. Fonseca, M. Egenhofer, C.A. Davis, and G. Câmara, "Semantic granularity in ontology-driven geographic information systems," *Annals of Mathematics and Artificial Intelligence*, vol. 36, pp. 121–151, 2002.

[3] D. Palacio, C. Derungs, and R. Purves, "Development and evaluation of a geographic information retrieval system using fine grained toponyms," *Journal of Spatial Information Science JoSIS*, vol. 11, pp. 1–29, 2015.

[4] A. Henrich and V. Lüdecke, "Characteristics of geographic information needs," in *Proc. of GIR'07*, Lisboa, Portugal, 2007, pp. 1–6.

[5] A. Ballatore, W. Kuhn, M. Hegarty, and E. Parsons, "Special issue introduction: Spatial approaches to information search," *Spatial Cognition & Computation*, vol. 16, no. 4, pp. 245–254, 2016.

[6] R. Jones, B. Rey, O. Madani, and W. Greiner, "Generating query substitutions," in *Proc. of WWW'2006*, Edinburgh, 2006, pp. 387–396.

[7] F. Fernández-Reyes, J. Hermosillo-Valadez, and M. Montes-y-Gmez, "A prospect-guided global query expansion strategy using word embeddings," *Information Processing & Management*, vol. 54, no. 1, pp. 1 – 13, 2018.

[8] C.-K. Huang, L.-F. Chien, and Y.-J. Oyang, "Relevant term suggestion in interactive web search based on contextual in formation in query session logs," *Journal of the American Society for Information Science and Technology*, vol. 54, no. 7, pp. 638–649, 2003.

[9] Y. Chen, G.-R. Xue, and Y. Yu, "Advertising keyword suggestion based on concept hierarchy," in *Proc. of WSDM '08*. New York, NY, USA: ACM, 2008, pp. 251–260.

[10] Y. Gupta and A. Saini, "A novel fuzzy-pso term weighting automatic query expansion approach using combined semantic filtering," *Knowledge-Based Systems*, vol. 136, pp. 97 – 120, 2017.

[11] H. Cao, D. Jiang, J. Pei, Q. He, Z. Liao, E. Chen, and H. Li, "Context-aware query suggestion by mining click-through and session data," in *Proc. of KDD '08*. New York, NY, USA: ACM, 2008, pp. 875–883.

[12] D. Jiang, J. Vosecky, K. W.-T. Leung, L. Yang, and W. Ng, "SG-WSTD: A framework for scalable geographic web search topic discovery," *Knowledge-Based Systems*, vol. 84, pp. 18 – 33, 2015.

[13] R. Laurini, "Geographic ontologies, gazetteers and multilingualism," *Future Internet*, vol. 7, pp. 1–23, 2015.

[14] K. Janowicz, S. Scheider, T. Pehle, and G. Ha, "Linkedgeodata: A core for a web of spatial open data," *Semantic Web Interoperability, Usability, Applicability*, vol. 3, no. 4, pp. 333–354, 2012.

[15] F. Fonseca, M. Egenhofer, P. Agouris, and G. Câmara, "Using ontologies for geographic information systems," *Transactions in GIS*, vol. 3, pp. 231–257, 2002.

[16] M. Codescu, G. Horsinka, O. Kutz, T. Mossakowski, and R. Rau, "OSMonto – an ontology of openstreetmap tags," in *Proc. State of the map Europe (SOTM-EU) 2011*, Vienna, Austria, 2011, pp. 55–64.

[17] A. Ballatore, M. Bertolotto, and D. Wilson, "Geographic knowledge extraction and semantic similarity in OpenStreetMap," *Knowledge and Information Systems*, vol. 37, no. 1, pp. 68–81, 2013.

[18] A. Voghera, R. Crivello, L.Ardissono, M. Lucenteforte, A. Savoca, and L. La Riccia, "Production of spatial representations through collaborative mapping. An experiment," in *Proc. of INPUT 2016*, 2016, pp. 356–361.

[19] L. Di Rocco, "Semantic enhancement of volunteered geographic information," in *Proc. of AI*IA-DC 2016*. Aachen, Germany: CEUR-WS.org, 2016, pp. 28–33.

[20] N. Mauro and L. Ardissono, "Session-based suggestion of topics for exploratory search," in *Proc. of ACM IUI 2018*. Tokyo, Japan: ACM, 2018, pp. 341–352.

[21] S. Gregory, "Finding overlapping communities in networks by label propagation," *New Journal of Physics*, vol. 12, no. 103018, 2010.

[22] A. Moro, A. Raganato, and R. Navigli, "Entity Linking meets Word Sense Disambiguation: a Unified Approach," *Transactions of the Ass. for Computational Linguistics (TACL)*, vol. 2, pp. 231–244, 2014.

[23] C. Manning, M. Surdeanu, J. Bauer, J. Finkel, S. Bethard, and D. Mc-Closky, "The Stanford CoreNLP natural language processing toolkit," in *ACL System Demonstrations*, 2014, pp. 55–60.