Object Location Estimation in Domestic Environments through Internet Queries

Ramon Izquierdo-Cordova, Eduardo F. Morales, L. Enrique Sucar
Instituto Nacional de Astrofísica, Óptica y Electrónica
{izquierdocr, emorales, esucar}@inaoep.mx

Abstract

Searching for objects is a fundamental task for service robots. Without any prior knowledge about the object's location, the robot needs to estimate the most probable place where to find the object in order to speed up the search process. In this work we automatically estimate likely locations of desired objects by combining information from four sources: Google, DBPedia, ConceptNet, and Word2Vec. It is shown that using any single source is not reliable enough so we used an ensemble of sources and combined their results using three methods: Borda count, simple combination, and weighted combination. We experimentally evaluate the benefits of the approach using information from users over more than 70 objects and propose two measures of quality to evaluate the results. We also compared the approach against other methods. The method has been used in a service robot to find objects in real-life conditions.

1 Introduction

Obtaining reliable information from Internet is a complex problem [Alani *et al.*, 2003]. Various knowledge databases have been automatically built with the purpose of providing machines comprehension skills by giving them access to information that humans consider as common sense. An example of common sense knowledge is that the most likely place to find, let's say an apple, is the kitchen. In applications like service robots, it is important to have such equivalent information to find objects within a house. Internet has become the main source of knowledge so it is natural to try to find such information in the Web.

We propose methods to estimate the probability that an object is located in each room of a domestic environment using automatic Internet queries; combining information from Google, DBPedia, ConceptNet, and Word2Vec using three techniques: Borda count, simple combination, and weighted combination. We also proposed two metrics to evaluate methods that generate probability estimations of object locations. The proposed approach was evaluated experimentally with 70 objects in realistic environments and compared with other methdos in the literature.

2 Related Work

Automatic estimation of the probability of locating an object within a domestic environment has been proposed by several authors. In [Kunze *et al.*, 2012] objects locations are estimated by counting entries in the *Open Mind Indoor Common Sense* database (OMICS) containing both, objects and rooms. Next, probabilities are computed using the conditional probability of the object given the room. A similar method employing OMICS is described in [Elfring *et al.*, 2014] where Lidstone's law is used to compensate unseen object-room combinations. Objects and scenes that a robot knows are used to predict the existence of new objects [Kollar and Roy, 2009]. The probabilities of finding an object at each location in the environment is calculated from the co-occurrence of the searched object with existing objects, form labels of scenes containing objects in the *Flickr* database. The same database is used in [Chumtong *et al.*, 2014] in a similar way with the purpose of estimating object-room relationships.

[Samadi *et al.*, 2012] present a method to obtain the probability distribution of the location of an object from Internet queries. It takes the results of a Web search engine when asked to find pages that contain the object and each of the rooms on the environment. The first 20 pages returned in queries of each pair object-room are classified using a support vector machine that determines to which page belongs each room, and probabilities are obtained by normalization. [Hanheide *et al.*, 2011] and [Lorbach *et al.*, 2014] quantify the likelihood of object-room pairs by counting the number of hits returned by a Web search (*Bing*) based on images and normalizing quantities with the hits returned by a query asking to search only the room. Probability estimation is used with an inverse purpose in [Viswanathan *et al.*, 2011], where the objective is to recognize a room from the chances of containing certain objects in it. Probabilities are estimated by counting the existence of objects in different rooms based on images stored in the *LabelMe* database.

Previous approaches are either restricted to the objects and rooms included in a particular database or use a single source for their queries. Another limitation is that they do not provide a quantifiable analysis of the appropriateness of their results, and are normally tested on a very small set of objects and rooms. This work avoids these limitations by using an ensemble of relatively general sources and by providing two metrics for comparison against human users.

3 Sources and Methods

There are several possible available sources that can be used to estimate the most likely place to find an object. In this research we used Google, DBPedia, ConceptNet, and Word2Vec.

3.1 Google

When a search is performed, Google returns the number of pages where matches (hits) with the indicated text were found. An initial way to estimate probabilities with *Google* is by searching pages with the object and the different rooms in the environment. So, if the search *apple+kitchen* has more hits than *apple+bedroom*, we can think it is more likely to find the apple in the kitchen than in the bedroom. A disadvantage of considering only hit numbers is that values are not proportional to the number of pages associated with each room. A solution to this problem is proposed in [Cilibrasi and Vitanyi, 2007], where the *Normalized Google Distance* (NGD) for two terms x and y is described:

$$NGD\left(x,y\right)=\frac{\max\left\{\log f\left(x\right),\log f\left(y\right)\right\}-\log f\left(x,y\right)}{\log N-\min\left\{\log f\left(x\right),\log f\left(y\right)\right\}} \tag{1}$$

being $f(x)$ and $f(y)$ the number of pages matching the terms x and y respectively; $f(x,y)$ is the number of pages matching both terms together; and N is the size of the *Google Index*. Google states that it has exceeded 60 billion of indexed pages [Inside Press, 2015], although other estimates give lower numbers [de Kunder, 2006]. The author of NGD indicates that the higher the value of N the more stable is the distance, so we decided to use Google's estimate.

3.2 DBPedia

Through automated extraction techniques, *DBPedia* constructs semantic information from plain text stored in Wikipedia. Once structured, the information can be consulted using SPARQL[PrudHommeaux *et al.*, 2008] which is available inside the DBPedia API functions[1]. Values for the relationship between the object and the room can be constructed with this database counting the number of Wikipedia articles in which, an object and an room appear simultaneously. The SPARQL query used in this work to retrieve the number of articles in Wikipedia with words *spoon* and *kitchen* is shown below:

```
select count(*) where {
  quad map virtrdf:DefaultQuadMap {
     graph ?g {
          ?s1 ?s1textp ?o1 .
          ?o1 bif:contains
          ' ( spoon AND kitchen ) '
          option ( score ?sc ) .
     }
  }
}
```

[1]http://dbpedia.org/sparql or http://live.dbpedia.org/sparql

3.3 ConceptNet

It is a semantic network that contains a large number of facts that a computer should know about the world, especially to understand written texts [Liu and Singh, 2004]. It is constructed by nodes that represent words or short phrases of natural language, and labeled relations between them. Nodes in *ConceptNet* are called *terms*. Figure 1 shows an extract of the ontology for the term *car* and its relation to the term *drive*. It shows a weight value ($weight = 1.0$), which can be used to estimate the probability that the object is inside that room.

As a semantic network whose concepts are connected by many dimensions, ConceptNet can approximate a value of relation between two terms not as a simple function of the number of hops between nodes representing these terms, but also considering the number and importance (weight) of all the routes connecting the two nodes. It is this relationship function between the object and the room, which we have used through the ConceptNet API available online[2], where we substitute *object* and *room* for existing terms such as *spoon* and *kitchen*.

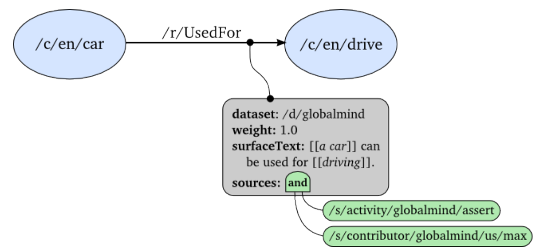

Figure 1: Part of ontology in *ConceptNet* describing the term *car* and its relation to *drive*

3.4 Word2Vec

The learning model *Word2Vec* [Goldberg and Levy, 2014] is defined as a set of algorithms that takes a corpus of text as input and generates a set of vectors representing the words in that corpus. It first builds a vocabulary with the analyzed text and then learn vectorial vocabulary representations using *Deep Learning* techniques. The resulting vectors can be used as features in natural language processing. In this paper we used pre-trained vectors from the *Google News* set containing about a hundred billion words. This model gets 300-dimensional vectors for three million words. Using the *cosine similarity* between the vectors representing the searched object and the room we can calculate a value for the object-room relationship. If we have two words represented by the vectors of dimension n, $A = \{A_1, ..., A_n\}$ and $B = \{B_1, ..., B_n\}$, the cosine similarity is computed by the formula:

$$similarity=\frac{\sum_{i=1}^{n}\left(A_i\times B_i\right)}{\sqrt{\sum_{i=1}^{n}A_i^2}\times\sqrt{\sum_{i=1}^{n}B_i^2}} \tag{2}$$

After obtaining the relative values for each object-room relationship according to the source used, a normalization process is followed so that the probability of all rooms sums 1.

[2]http://conceptnet5.media.mit.edu/data/5.2/assoc/c/en/object?filter=/c/en/room&limit=1

3.5 Combining Estimates from Different Sources

Estimating the probabilities of finding an object in different rooms with a single source presents two problems: (i) the source may not contain information about the objects we are searching for, and (ii) the quality of the results varies depending on the object being searched. For instance, for Conceptnet, it is more likely to find milk in the laundry room than in the kitchen, for Word2Vec, it is more likely to find the fabric softener in the living room than in the laundry room, for Google, it is more likely to find a laptop in the bathroom than in the bedroom, and for DBPedia, it is more likely to find a DVD in the patio than in the living-room.

Ensembles methods have been successfully used in machine learning to mitigate the errors from single classifiers. In this paper we used an ensemble of the four sources and evaluated three different ways to combine their results. Each source returns an ordered list of rooms with an associated probability. One of the easiest ways to make a combination of ordered lists is using the Borda Method [de Borda, 1781]. It only takes into account the position in the list (i.e., not the probability values). In this paper, we use it in two ways: for combining the results from the Internet sources and for merging the users' opinions that are used as a reference in the evaluation process as explained in Section 4.1.

Given N candidates and multiple voters where each voter arranges candidates in descending order according to preferences, $N-1, N-2, ..., 0$ points are assigned to the first, second and last candidate respectively as selected by each voter. The points of each candidate are added up for all voters and the winning candidate is the one with most points. If we think that each source is a voter and if r_{ik} is the position of the room i in the ordered list produced by the information from source k, then the Borda count for room i is given by the formula:

$$b_i = \sum_{k=1}^{N} (N - r_{ik}) \tag{3}$$

It has been proved that this method is an optimal positional voting according to various standards [Lansdowne and Woodward, 1996]. In particular, it determines the consensus that minimizes the sum of squared deviations of the selection of each voter and reduces the number and class of voting paradoxes [Cook and Seiford, 1982].

Another way to generate a combination of results is to take the average of the probability estimates generated by each source. This allows to take into account the differences in probability values between the rooms. If a source assigns a high probability to a room, this may be enough to reach one of the first places in the list even if the other sources produce low probability values. The simplest combination P_S for estimating probabilities is taking the average:

$$P_S = (P_G + P_D + P_C + P_W)/4 \tag{4}$$

being P_G, P_D, P_C y P_W the probability estimates obtained from *Google*, *DBPedia*, *ConceptNet* and *Word2Vec*, respectively.

Different sources may have different degrees of precision so a more principled way of combining them is to estimate their reliability and use it when combining their results. We estimated the reliability of each source from half of the objects used in our experimental set and used a weighted combination of sources. The probability value P_W of the ensemble is given by the formula:

$$P_W = W_G P_G + W_G P_D + W_G P_C + W_W P_W \tag{5}$$

where the W_i's are the estimates for each source.

Once the combination values are obtained for each room they can be normalized.

4 Evaluation

We need a reference value in order to evaluate the proposed methods. We also would like the method to be used in realistic conditions. With that purposes, we conducted a survey to seven people who were asked to list ten small objects existing in their homes. Additionally we include two sets of objects from competitions in robotics and computer vision (*Semantic Robot Vision Challenge 2009* and *ImageCLEF 2014*). At the end we used 72 objects as shown in Table 1.

We considered an environment conformed by ten rooms shown in Table 2. We asked the same group of people to answer a second survey where they decided in which order should the set of ten rooms be explored in order to find each of the 72 objects. The survey is shown in Figure 2.

Kitchen	Bedroom	Bathroom	Living room
Dining room	Study	Play room	Patio
Laundry room	Garage		

Table 2: Existing rooms in our tested environment.

User opinions from the second survey were combined using the Borda count method. Considering the most likely place, according to the users' average opinion, most of the objects are located in the kitchen and the bedroom as shown in Figure 3.

4.1 Evaluation Metrics

We have an ordered list of rooms from the users and we want to compare it to ordered lists with a probability associated

Figure 2: Online survey to consult users on the room exploration order for finding objects (interface in Spanish).

ac remote	bookshelf	coat	dvd	hand cream	milk	potato chips	softener
apple	bottle	coffee	extinguisher	handbag	mop	potty	spoon
baby bottle	bread	coke	fly swatter	headphone	nail clipper	printer	tablet
backpack	broom	comb	food	inhaler	orange	pumpkin	towel
ball	cd	cookie	fork	jacket	paddle	scissor	toy
bed sheet	cell charger	cosmetic bag	fridge	key	pen	shirt	toy car
bible	cellphone	cracker	frying pan	knife	phone	shoe	trash
blanket	chair	cup	glass of water	laptop	pillow	soap	tv remote
book	clothe	dinosaur	glasses	medicine	plate	sock	videogame

Table 1: Objects used in the probability estimation experiments obtained by consulting 15 users.

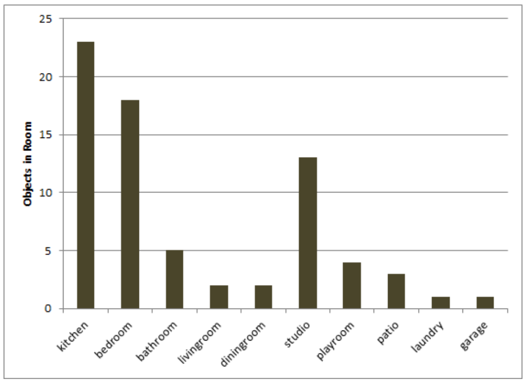

Figure 3: Distribution inside the rooms of the objects used to estimate the probability, according to the opinions of all users computed from Borda count.

obtained from the different sources. We can either ignore the probability values and make a comparison between two ordered lists or we can associate a probability to the user's list and compare two probability distributions. We followed both approaches.

A measure to compare ordered lists is the Kendall's Tau distance [Kendall, 1938]. Let $[n] = \{1, ..., n\}$ be a universe of elements and S_n the set of all permutations of $[n]$. For $\sigma \in S_n$ let $\sigma(i)$ the position within the list of element i, then the Kendall Tau distance $K(\sigma)$ is defined by:

$$K(\sigma) = \sum_{(i,j):i>j} [\sigma(i) < \sigma(j)] \quad (6)$$

which measures the total inversions between pairs of elements. The main disadvantage of the Kendall's Tau distance is that it does not take into account the position of the elements even though they may have different importance. This is why authors ([Shieh, 1998], [Kumar and Vassilvitskii, 2010]) have proposed modifications to this metric to satisfy the requirements of importance. There are many ways to define a weighted analog for the Kendall's Tau distance. A summary of weighted measurements is given in [Webber *et al.*, 2010]. In particular, in this paper we used a penalty for the inversion of the elements i and j as the product of their weights $w_i w_j$ [Kumar and Vassilvitskii, 2010], so the value of weighted Kendall's Tau distance $K(\sigma)$ is given by:

$$K_w(\sigma) = \sum_{(i,j):i>j} w_i w_j [\sigma(i) < \sigma(j)] \quad (7)$$

where for each position $i \in \{1 \ldots n\}$ in the list, we define $w_i > 0$ as the weight of element i, and $w = \{w_1, ..., w_n\}$ as the set of weights.

The weight for the different sources are given by their estimated probabilities. In order to assign weights to the users, we first counted how many times each user matches the vote of all the users obtained with the Borda method. The results are shown in Figure 4. This figure shows that on average, users agree on the most likely room over 80% of the time while for rooms that are in the middle of the list (positions 5, 6 and 7), the percentage of agreement is about 25%. We can also notice that humans tend to agree as well on where not to find a particular object. This suggests a very high probability of finding an object in the first rooms suggested by the users and a very low probability of finding the object in the last rooms. A reasonable probability distribution for this behavior is an exponential function. So for assigning the weights of the users for the Kendall Tau distance, we used an exponential function ($e^{-\lambda x}$) that fits the coincidence of the users on the first rooms, with $\lambda = 2.0$, where $x = 0, 1, ..$ is the position in the list being $x = 0$ the first position.

We also compared every position on the list of the proposed methods against the Borda count of the users. This is shown in Figure 5. We can see that the proposed methods tend to coincide in the first rooms but show little coincidences in the last rooms. This is probably not surprising as the information stored in the different sources is mainly about related objects and not about where not to find an object. From this figure,

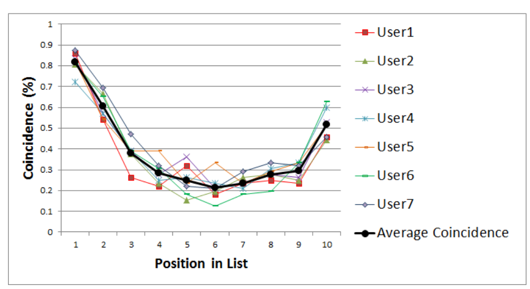

Figure 4: Coincidence of user opinions against Borda voting of all users measured by room position in the exploration list. The picture shows the average distances to 72 objects.

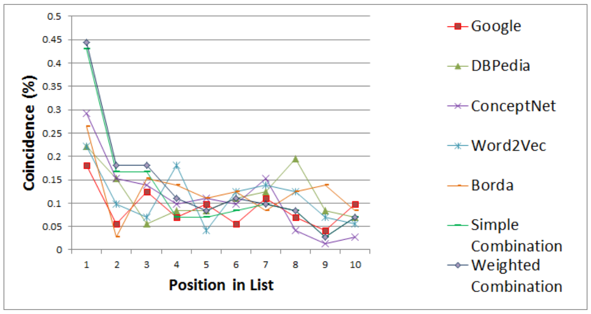

Figure 5: Coincidence of proposed methods against Borda voting of all users measured by room position in the exploration list. The picture shows the average distances to 72 objects.

we can clearly appreciate the advantages of using an ensemble of sources as they obtained close to 45% of coincidences in the first room (out of 10) while the sources percentages, when considered in isolation, range between 18% (Google) to 29% (ConceptNet). On the other hand, if we consider coincidences of the first room within the first 3 and 5 rooms selected by humans, we obtained a coincidence of 65% and 77% respectively. In service robot applications it is important to search in a small number of rooms. With the proposed approach the probability of finding an object within the first three rooms (out of 10), according to a set of users, is 65%.

Figure 6 shows the results of the individual sources used as well as their ensembles using as reference the Borda vote of all users. Measurements are taken with Kendall's Tau distance in its weighted version. Under this measure, the best results are those obtained with the ensembles.

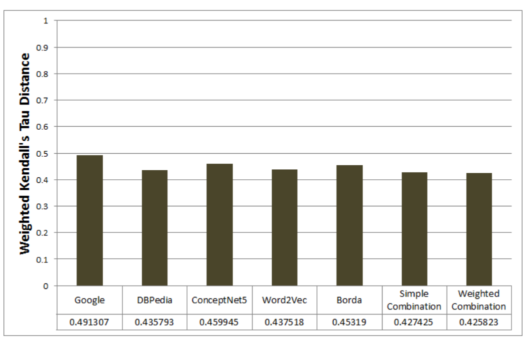

Figure 6: Weighted Kendall's Tau distance of the proposed methods for probability estimation of the location of the 72 objects (average) in the rooms of a home environment measured against Borda voting of all users. Values close to 0 indicate greater similarity between lists.

For comparison we show the weighted Kendall's Tau distance obtained from the evaluation of the opinion of individual users participating in the survey (see Figure 7). Since the Borda count of the user is used as reference, it is not too surprising that these results are better that the proposed methods. It also shows certain consistency among users which is consistent with the results shown in Figure 5 where there are high

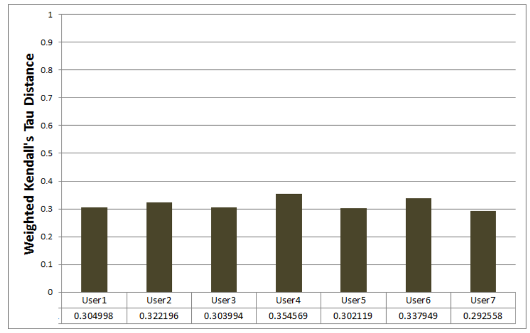

Figure 7: Weighted Kendall's Tau distance of individual opinions of the location of the 72 objects (average) in the rooms of a home environment measured against Borda voting of all users. Values close to 0 indicate greater similarity between lists.

Method	Hellinder
Google	0.1957
DBPedia	0.5943
ConceptNet	0.5632
Word2Vec	0.2217
Borda	0.2647
SEns	0.1916
WEns	0.1916

Table 3: Hellinger distance between the exponential distribution following the Borda count of the users and the different sources, where *SEns* is the ensemble with a average of the probabilities and *WEns* is the ensemble with the weighted combination of sources.

levels of coincidences.

Another way to compare the results from the users against the estimates from the different sources is assigning a probability value to the users. Using an exponential function for the users, we compare it against the probability distributions of the individual sources and the ensembles using the Hellinger distance[3]. For discrete probability distributions $P = p_1, p_2, \ldots, p_k$ and $Q = q_1, q_2, \ldots q_l$ it is defined as:

$$H(P,Q) = \frac{1}{\sqrt{2}}\sqrt{\sum_{i=1}^{k}(\sqrt{p_i} - \sqrt{q_i})^2}$$

Table 3 shows the results. It can be seen that in general the ensembles are closer to the probability distribution of the Borda count of the users assuming an exponential distribution.

Making a comparison against the reported results in the literature is not easy as some of sources used to obtain the probability values are not longer available, and the objects and type of rooms are all different. We present a partial, qualitative comparison with other three approaches, by contrasting the probabilities they obtain for some objects/rooms against

[3]Other measures, like Kullback-Leibler, do not apply as some probabilities for certain objects and rooms are zero.

	Article results				Weighted Combination			
	cup	glass	sandwich	chair	cup	glass	sandwich	chair
auditorium	0.070	0.075	0.085	0.050	0.059	0.069	0.046	0.099
classroom	0.070	0.075	0.085	0.160	0.052	0.061	0.078	0.121
conference room	0.070	0.075	0.085	**0.270**	0.107	0.120	0.060	0.118
hallway	0.060	0.065	0.075	0.040	0.060	0.128	0.054	0.103
kitchen	**0.240**	0.140	0.220	0.060	**0.240**	**0.312**	0.215	0.143
laboratory	0.070	**0.230**	0.085	0.050	0.059	0.067	0.058	0.090
locker room	0.070	0.075	0.085	0.050	0.206	0.102	0.109	0.092
office	0.135	0.040	0.045	0.160	0.072	0.064	0.062	**0.149**
restaurant	0.210	0.220	**0.250**	0.150	0.145	0.076	**0.319**	0.083

Table 4: Results reported in *Searching Objects in Large-scale Indoor Environments: A Decision-theoretic Approach*, Lars Kunze, Michael Beetz, Manabu Saito, Haseru Azuma, Kei Okada, Masayuki Inaba [Kunze *et al.*, 2012]. The results are presented in the article in a graph, rather than a table, so the numbers are approximated.

Article results

	laptop	papers	cup	coffee
bathroom	0.070	0.100	**0.360**	0.220
printer room	0.230	**0.570**	0.160	0.210
kitchen	0.130	0.120	0.290	**0.300**
office	**0.560**	0.220	0.180	0.280

Weighted Combination

	laptop	papers	cup	coffee
bathroom	0.149	0.275	0.193	0.318
printer room	**0.389**	0.204	0.325	0.126
kitchen	0.213	0.200	**0.359**	**0.392**
office	0.249	**0.321**	0.122	0.165

Table 5: Results reported in *Enabling Robots to Find and Fetch Objects by Querying the Web* (Extended Abstract). Thomas Kollar, Mehdi Samadi, Manuela Veloso [Kollar *et al.*, 2012].

Article results

	book	glass	knife	milk	news-paper	shirt
bathroom	0.300	**0.600**	0.200	0.200	0.300	**0.700**
bedroom	0.200	0.300	0.100	0.100	0.100	0.300
kitchen	**0.400**	0.400	**0.400**	**0.600**	**0.400**	0.600
living	0.300	0.400	0.100	0.200	0.300	0.100

Weighted Combination

	book	glass	knife	milk	news-paper	shirt
bathroom	0.304	0.243	0.187	0.234	**0.303**	0.314
bedroom	**0.343**	0.212	0.179	0.210	0.270	**0.428**
kitchen	0.25	**0.412**	**0.523**	**0.370**	0.279	0.166
living	0.101	0.132	0.110	0.185	0.147	0.092

Table 6: Results reported in *Prior-Assisted Propagation of Spatial Information for Object Search* Malte Lorbach, Sebastian Höfer, Oliver Brock [Lorbach *et al.*, 2014]. The results are presented in the article without normalization.

our ensamble method. The results are shown in Tables 4, 5, and 6.

Although the analysis of the results is subjective and could depend on each particular environment, we can appreciate that in general the proposed method produces similar or better results than the other approaches. In Table 4, both methods produce reasonable results for all the objects. Table 5 shows common sense results for both methods for three of the objects, but for *cup* the *kitchen* seems more probable than *bathroom*. The results in Table 6 seem better for our approach in at least three objects: *book, glass, shirt*. An important future challenge for the service robots community will be to generate a benchmark in this area so that the different techniques can be compared.

Our method is not restricted to any particular set of objects or rooms and combines information from more than one source which, as shown in our experiments, produces better results than using any single source.

Since the proposed method returns an ordered list of rooms in less than a second using a standard computer, we used it online with a real robot. In the experiments, the user specified an object and the robot searched in its current map of the environment finding the object in the first or second in an environment with 6 rooms.

5 Conclusions

We have proposed a mechanism to automatically estimate the probabilities of locating an object in different rooms of a domestic environment using a weighted combination of information gathered from four different sources. Contrary to previous work, we have also proposed two metrics to evaluate methods that generate probability estimations of object locations. An evaluation is also made comparing results against the opinion of a group of users employing our measure of similarity based on the weighted Kendall's Tau distance over a large number of objects. We have also qualitatively compared our method against other reported work. In general, our results are very competitive against other methods and we were able to estimate the performance of the approach against a set of users. We have also used our method on a service robot under realistic conditions with promising results. As future work, we are exploring other ways to improve the probability estimates and include more powerful searching and recognition methods of objects for a real robot.

References

[Alani *et al.*, 2003] Harith Alani, Sanghee Kim, David E Millard, Mark J Weal, Wendy Hall, Paul H Lewis, and Nigel R Shadbolt. Automatic ontology-based knowledge extraction from web documents. *Intelligent Systems, IEEE*, 18(1):14–21, 2003.

[Chumtong *et al.*, 2014] Puwanan Chumtong, Yasushi Mae, Kenichi Ohara, Tomohito Takubo, and Tatsuo Arai. Object search using object co-occurrence relations derived from web content mining. *Intelligent Service Robotics*, 7:1–13, 2014.

[Cilibrasi and Vitanyi, 2007] Rudi L Cilibrasi and Paul MB Vitanyi. The google similarity distance. *IEEE Transactions on Knowledge and Data Engineering*, 19(3):370–383, 2007.

[Cook and Seiford, 1982] Wade D Cook and Lawrence M Seiford. On the Borda-Kendall consensus method for priority ranking problems. *Management Science*, 28(6):621–637, 1982.

[de Borda, 1781] Jean C de Borda. Mémoire sur les élections au scrutin. *Mémoire de l'Academie Royale des Sciences de Paris*, pages 657–664, 1781.

[de Kunder, 2006] Maurice de Kunder. The size of the World Wide Web, 2006. Accessed 2015-05-03.

[Elfring *et al.*, 2014] Jos Elfring, Simon Jansen, René van de Molengraft, and Maarten Steinbuch. Active object search exploiting probabilistic object–object relations. In *RoboCup 2013: Robot World Cup XVII*, pages 13–24. Springer, 2014.

[Goldberg and Levy, 2014] Yoav Goldberg and Omer Levy. Word2Vec explained: deriving Mikolov et al.'s negative-sampling word-embedding method. *arXiv preprint arXiv:1402.3722*, 2014.

[Hanheide *et al.*, 2011] Marc Hanheide, Charles Gretton, Richard Dearden, Nick Hawes, Jeremy Wyatt, Andrzej Pronobis, Alper Aydemir, Moritz Göbelbecker, and Hendrik Zender. Exploiting probabilistic knowledge under uncertain sensing for efficient robot behaviour. In *IJCAI Proceedings-International Joint Conference on Artificial Intelligence*, volume 22, page 2442, 2011.

[Inside Press, 2015] Google Inside Press. Google inside search, 2015. Accessed 2015-05-21.

[Kendall, 1938] Maurice G Kendall. A new measure of rank correlation. *Biometrika*, pages 81–93, 1938.

[Kollar and Roy, 2009] Thomas Kollar and Nicholas Roy. Utilizing object-object and object-scene context when planning to find things. In *IEEE International Conference on Robotics and Automation. ICRA'09*, pages 2168–2173. IEEE, 2009.

[Kollar *et al.*, 2012] Thomas Kollar, Mehdi Samadi, and Manuela Veloso. Enabling robots to find and fetch objects by querying the web. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems*, volume 3, pages 1217–1218. International Foundation for Autonomous Agents and Multiagent Systems, 2012.

[Kumar and Vassilvitskii, 2010] Ravi Kumar and Sergei Vassilvitskii. Generalized distances between rankings. In *Proceedings of the 19th international conference on World Wide Web*, pages 571–580. ACM, 2010.

[Kunze *et al.*, 2012] Lars Kunze, Michael Beetz, Manabu Saito, Haseru Azuma, Kei Okada, and Masayuki Inaba. Searching objects in large-scale indoor environments: A decision-theoretic approach. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4385–4390. IEEE, 2012.

[Lansdowne and Woodward, 1996] Zachary F Lansdowne and Beverly S Woodward. Applying the Borda ranking method. *Air Force Journal of Logistics*, 20(2):27–29, 1996.

[Liu and Singh, 2004] Hugo Liu and Push Singh. Conceptnet - a practical commonsense reasoning tool-kit. *BT technology journal*, 22(4):211–226, 2004.

[Lorbach *et al.*, 2014] Malte Lorbach, Sebastian Hofer, and Oliver Brock. Prior-assisted propagation of spatial information for object search. In *International Conference on Intelligent Robots and Systems (IROS 2014)*, pages 2904–2909. IEEE, 2014.

[PrudHommeaux *et al.*, 2008] Eric PrudHommeaux, Andy Seaborne, et al. Sparql query language for rdf. Technical report, W3C, 2008.

[Samadi *et al.*, 2012] M. Samadi, T. Kollar, and M. Veloso. Using the web to interactively learn to find objects. In *Proceedings of the 26th Conference on Artificial Intelligence*, pages 2074–2080. Association for the Advancement of Artificial Intelligence, 2012.

[Shieh, 1998] Grace S Shieh. A weighted kendall's tau statistic. *Statistics & Probability Letters*, 39(1):17–24, 1998.

[Viswanathan *et al.*, 2011] Pooja Viswanathan, Tristram Southey, James Little, and Alan Mackworth. Place classification using visual object categorization and global information. In *Canadian Conference on Computer and Robot Vision (CRV)*, pages 1–7. IEEE, 2011.

[Webber *et al.*, 2010] William Webber, Alistair Moffat, and Justin Zobel. A similarity measure for indefinite rankings. *ACM Transactions on Information Systems (TOIS)*, 28(4):20, 2010.