Category-Driven Approach for Local Related Business Recommendations

Yonathan Perez [*]
Stanford University
yperez@cs.stanford.edu

Michael Schueppert
Google New York
mschueppert@google.com

Matthew Lawlor
Google New York
mflawlor@google.com

Shaunak Kishore
Google New York
skishore@google.com

ABSTRACT

When users search online for a business, the search engine may present them with a list of related business recommendations. We address the problem of constructing a useful and diverse list of such recommendations that would include an optimal combination of substitutes and complements. Substitutes are similar potential alternatives to the searched business, whereas complements are local businesses that can offer a more comprehensive and better rounded experience for a user visiting the searched locality. In our problem setting, each business belongs to a category in an *ontology of business categories*. Two businesses are defined as substitutes of one another if they belong to the same category, and as complements if they are otherwise relevant to each other.

We empirically demonstrate that the related business recommendation lists generated by Google's search engine are too homogeneous, and overemphasize substitutes. We then use various data sources such as crowdsourcing, mobile maps directions queries, and the existing Google's related business graph to mine association rules to determine to which extent do categories complement each other, and establish relevance between businesses, using both category-level and individual business-level information. We provide an algorithmic approach that incorporates these signals to produce a list of recommended businesses that balances pairwise business relevance with overall diversity of the list. Finally, we use human raters to evaluate our system, and show that it significantly improves on the current Google system in usefulness of the generated recommendation lists.

1. INTRODUCTION

Related business recommendations help users discover businesses or establishments that would be useful when searching for a specific source business. Search engines today answer this need by providing recommendations similar to the searched business, relying mostly on search-based association rules of the form "users who search for business x also search for business y". Such recommendations are often *substitutes* to the searched business - establishments that the user could visit instead of the searched business.

Yet, such substitutes may not be the only relevant recommendations for a source business. In some cases a user could benefit more from recommended *complements* – local businesses that one could visit in conjunction to the source business and that could provide the user with a better rounded experience at the searched locality. For example, a user searching for a *specific* hotel might benefit more from recommendations for good restaurants and attractions near that hotel than from only recommendations for other hotels. Some major commercial recommender systems use both substitute and complement notions of relevance in making product recommendations [7, 10]. We aim to leverage both notions of relevance in generating a list of recommended local businesses given a source business.

The usefulness of recommended complements for a source business x can depend on many factors. Later in the paper we show that the type, or category, of x affects heavily on the usefulness of complementary recommendations for it, as well as on the categories of other businesses that could serve as useful complements.

Another important factor is the user's intent - whether she is considering x among other alternatives, or she is planning a visit to x or even already at the locality of x. User intent was not modeled in this paper, and is a topic for future work. However, our approach of combining complements in the related business recommendations, rather than only similar substitutes, is geared towards users who plan a visit to the source business. Thus, lessons from the approach taken in this paper could be more relevant in settings where such user intent is more likely (e.g. mobile searches where the user is near the searched locality, or on the way there).

Our problem setting is that of a non-personalized recommender system, and is demonstrated in Figure 1. When an anonymous user searches for a specific business name, in addition to search results for that business, the system provides a list of k related business recommendations. We aim to improve user satisfaction from the generated recommendation list by constructing a relevant and diverse list of substitute and complement recommendations using both category and business-specific information.

[*] This work was mostly done while at Google New York

CIKM'15, October 19–23, 2015, Melbourne, Australia.
ACM 978-1-4503-3794-6/15/10.
DOI: http://dx.doi.org/10.1145/2806416.2806495 .

In our problem setting, there exists an ontology of business categories. Each business is a member of some subset of categories in this ontology. In our approach, the relevance of a potential related business recommendation y to the source business x can be established by looking at: (i) the relevance of the set of categories to which y belongs to the set of categories to which x belongs, and (ii) the relevance of y to x given a recommendation category. If the recommendation category is common to both x and y, then y is a substitute recommendation for x. Otherwise, it is a complement recommendation.

One simple approach is to recommend the k businesses that are most relevant to the searched business. Yet, such an approach would typically yield a very homogeneous recommended business list, with the majority of businesses belonging to the category c most relevant to the searched business - i.e. a list of only complements, or more commonly - only substitutes as the source business itself often belongs to c. While each recommended business is highly relevant in these cases, the list as a whole can be redundant. Thus, it is desirable to recommend relevant businesses of varying types and include both substitutes and various types of complements. We use the ontology defined over the business categories to define and promote intra-list diversity, by penalizing recommendations of multiple businesses of the same category. We explore and compare several methods for such penalization.

Unfortunately, as described later in section 4, such penalization in settings where every business can belong to multiple categories makes the problem of finding an optimal recommendation list NP-hard. To make the problem solvable in polynomial time, we associate every business with a single category of interest.

We use crowdsourcing and association rules mined from category-aggregated mobile maps directions queries to establish category-to-category relevance scores. To determine relevance between businesses given category information, we use current Google's business recommendation scores as a baseline. We then adjust these scores using the directions queries data. Using these signals, we establish business-to-business relevance scores, and we produce top-k recommendation lists that optimize for these scores along with category redundancy penalties.

The main contributions of our paper are:

(1) An empirical demonstration that users find complementary related business recommendations useful, whereas the current Google system is strongly focused on making recommendations similar to the searched business. This evidence is shown clearly in Figure 2 and discussed in section 3.

(2) An algorithmic approach for constructing a useful, and often diverse, related business recommendation list in a setting where a business categorization exists. The constructed list consists of a useful mixture of substitute and complement recommendations. This approach is presented in section 4.

(3) Establishment of relevance scores between businesses that take into account the business category ontology. We used multiple data sources to mine and construct the relevance scores: crowdsourcing, mobile maps directions queries, and the current Google related business graph. The construction of relevance scores is explained in section 5.

(4) Large-scale human-based empirical evaluation of several variations of our system. We demonstrate that our system achieves a significant improvement over the current Google system in recommendation list usefulness. The evaluation process and the results and insights that arise from it are discussed in section 6.

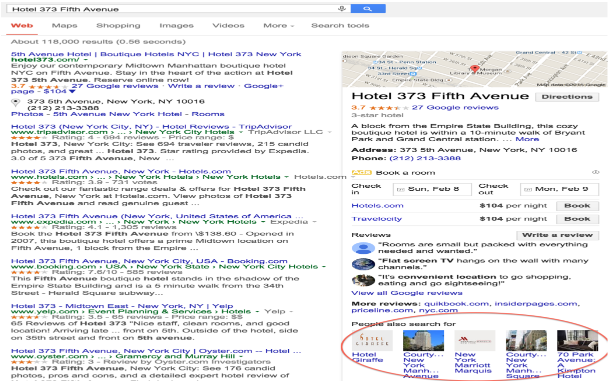

Hotel 373 Fifth Avenue

About 118,000 results (0.56 seconds)

5th Avenue Hotel | Boutique Hotels NYC | Hotel 373 New York
hotel373.com/
Enjoy our contemporary Midtown Manhattan boutique hotel NYC on Fifth Avenue. Stay in the heart of the action at **Hotel 373 5th Avenue**. Reserve online now!
3.7 ★★★★☆ 27 Google reviews · Write a review · Google+ page - $104
373 5th Avenue, New York, NY 10016
(212) 213-3388
Photos - 5th Avenue New York Hotel - Rooms

Hotel 373 (New York City, NY) - Hotel Reviews - TripAdvisor
www.tripadvisor.com › ... › New York City Hotels - TripAdvisor LLC
★★★★☆ Rating: 4 - 694 reviews - Price range: $
Hotel 373, New York City: See 694 traveler reviews, 215 candid photos, and great ... **Hotel 373**. Star rating provided by Expedia. 3.0 of 5 373 **Fifth Avenue**, New ...

Hotel 373 Fifth Avenue, New York - Hotels.com
www.hotels.com › ... › New York Hotels › New York Hotels - Hotels.com
★★★★☆ Rating: 3.9 - 731 votes
Check out our fantastic range deals & offers for **Hotel 373 Fifth Avenue**, New York at Hotels.com. View photos of **Hotel 373 Fifth Avenue** and read genuine guest ...

Hotel 373 Fifth Avenue (New York, United States of America ...
www.expedia.com › ... › New York › New York Hotels - Expedia
★★★★☆ Rating: 4.1 - 1,305 reviews
Book the **Hotel 373 Fifth Avenue** from \$138.60 - Opened in 2007, this boutique hotel offers a prime Midtown location on Fifth Avenue, 1 block from the Empire ...

Hotel 373 Fifth Avenue, New York City, USA - Booking.com
www.booking.com › USA › New York State › New York City Hotels
★★★★☆ Rating: 7.6/10 - 585 reviews
This **Fifth Avenue** boutique **hotel** stands in the shadow of the Empire State Building and is a 5 minute walk from the 34th Street - Herald Square subway...

Hotel 373 - Midtown East - New York, NY | Yelp
www.yelp.com › Event Planning & Services › Hotels - Yelp
★★★★☆ Rating: 3.5 - 65 reviews - Price range: $$
65 Reviews of **Hotel 373** "Nice staff, clean rooms, and good location! Arriving late ... front on 5th. Outside of the hotel, side on 35th street and front on **5th avenue**.

Hotel 373 Fifth Avenue, New York City | Oyster.com -- Hotel ...
www.oyster.com › ... › Gramercy and Murray Hill
★★★☆☆ Rating: 3 - Review by Oyster.com Investigators
Hotel 373 Fifth Avenue, New York City: See 176 candid photos, pros and cons, and a detailed expert hotel review of

Hotel 373 Fifth Avenue
3.7 ★★★★☆ 27 Google reviews
3-star hotel
A block from the Empire State Building, this cozy boutique hotel is within a 10-minute walk of Bryant Park and Grand Central station. ... More
Address: 373 5th Avenue, New York, NY 10016
Phone: (212) 213-3388

Ads Book a room

Check in	Sun, Feb 8	Check out	Mon, Feb 9
Hotels.com		**$104** per night	Book
Travelocity		**$104** per night	Book

Reviews
"Rooms are small but packed with everything needed and wanted."
"**Flat screen TV** hangs on the wall with many channels."
"It's **convenient location** to go shopping, eating and go sightseeing!"
View all Google reviews
More reviews: quikbook.com, insiderpages.com, priceline.com, nyc.com

People also search for
Hotel Giraffe
Courty... New York Manh... Avenue
New York Marriott Marquis
Courty... New York Manh... Square
70 Park Avenue: A Kimpton Hotel

Figure 1: The related business recommendations (circled in red) for the source business ”Hotel 373 Fifth Avenue” - a boutique hotel in Midtown East, Manhattan. All of the recommendations by the existing system are similar hotels, that serve as substitutes for the source business.

2. RELATED WORK

Item-to-item relevance notions of substitutes and complements have been used by some recommender systems [7, 10, 21]. In these personalized product recommender systems, substitute and complement relations were mined from user browsing and purchasing data. In such e-commerce product recommender systems, different notions of relevance can be used for different stages of the user experience - i.e. recommending substitutes before the user makes a purchase and complements after the purchase. Related business recommendations differ in the sense that its harder for the system to acquire information about different stages in the user behavior and clear signals of when to recommend substitutes and when to recommend complements. Thus, we took the approach of constructing a recommendation list that mixes the two.

Business categorization in our system serves a double purpose – on one hand, we use category information as a factor in determining business-to-business relevance by measuring relevance of category sets to each other, and on the other hand, we use categories to limit redundancy in the produced recommendation list. Thus, in our use of category information to construct a mix of substitutes and complements to the source business, we use *implicit diversification* of the recommendation list as a mean to reduce redundancy and increase user satisfaction. The settings of our problem are those of a non-personalized recommender system, and share traits of both traditional recommender system settings and web search settings. On one hand, the system discovers new useful recommendations that were not searched by the user, and on the other hand, the system is non-personalized and provides its suggestions in response to a web search query for a specific business.

Significant work has been done on result diversification in both recommender systems [18, 14, 20, 22, 13] as well as web search [6, 8, 11, 12, 19, 5]. There are two main notions of such diversification in literature [4] – one is resolving ambiguity and reaching high or sufficient coverage of several topics [2, 12, 14, 19, 5], and the other is avoiding redundancy [22, 3, 20]. A recent work, [13], is also aiming to address both notions.

Multiple approaches have been suggested in coping with ambiguous queries or covering multiple meanings or subtopics. One such approach is the IA-Select [2], where the settings include a taxonomy of the results and a distribution of user intents over this taxonomy given an ambiguous query. IA-Select then aims to maximize the probability that at least one returned result is relevant to the user. Another well-known approach is xQuAD [12], where ranking results for an ambiguous query is done by estimating how well does a result cover an uncovered aspect of the answer. In our problem settings however, coverage is not an objective. We do not aim to cover any number of categories, and we do not try to guarantee at least one useful result, but to maximize user satisfaction from the entire recommendation list, hoping that the user would find as many as possible of the provided recommendations useful. In [5], the goal is to return a result list with per-topic coverage that is proportional to that topic's popularity. While the number of recommendations from a specific category is often higher for categories that are highly relevant to the source business, this is not an objective of our system. Categories assist in determining relevance of businesses to the source business, but are not used for impose any constraints on the composition or coverage of the generated recommendation list.

Our system does aim to reduce redundancy of multiple businesses from the same category and implicitly increase intra-list diversity. In [22], intra-list diversification was pursued as a mean for user satisfaction from book recommendations in a personalized collaborative-filtering system. Similarity and diversity were measured using a taxonomy of books, and the intra-list similarity between items was minimized using a greedy algorithm. The output of the system was combined with a non-diversified list using a diversification factor. Our approach differs significantly in the sense that we do not scale a "diversified list" with an "accurate list", but introduce diversification natively into our optimization objective by enforcing a diminishing returns effect of multiple recommendations from the same category. Explicitly penalizing redundancy was done in [3]. Instead of ranking results by similarity to the query, the MMR (Maximal Marginal Relevance) method uses a linear combination of that similarity with a penalty for similarity to previous results. However, MMR does not take into account a diminishing returns effect and does not increase penalty severity with increased redundancy.

Another aspect touched by our system is that of locality aware recommendations. This is done implicitly as scores of local businesses tend to be higher in the signals we use to establish business-to-business relevance scores. While location is only used implicitly in our system, location-awareness plays a more explicit part in several locality-aware recommender systems [9, 17]. Location is actively used in ranking results for hyper-local web queries [15], and indeed used in determining the relevance between businesses in Google related business graph (see next section).

3. BACKGROUND AND MOTIVATION

Google's current source of related business recommendations is the *Related Business Graph (RBG)*. The RBG is a directed graph where a business x is linked to several (typically hundreds) related businesses and the weight of an edge (x, y) expresses the relevance of y to x. There are multiple signals that contribute to such relevance scores, and we cannot disclose them in this paper. However, search-based association rules (i.e. "users who search x also search for y") are a major component in the RBG's edge weights. Additional important signals include the geographical proximity of y to x (closer businesses are preferred), as well as various measures of *similarity* between x and y. The current RBG-based system produces related business recommendations for a searched business x by returning the top k neighbors of x in the RBG by weight. The RBG-based system is the main system used by Google in production, and for most source business queries [1], the related business recommendation list displayed to the user is produced by that system. Such recommendations tend to be similar to the searched business and to one another, and mostly constitute of substitutes to the source business.

Yet, similar businesses are not always what users want to see as recommendations when they search for a business. To demonstrate this we turned to crowdsourcing and conducted a survey, asking users what information they would find useful when searching for a business from a specific category of interest. The survey was conducted using Google Consumer Surveys[2], and includes 372 participants (web users accessing premium content online).

In the survey, and throughout this paper, we limited ourselves to the geographical scope of New York City (NYC). This controls for geographic settings and cultures, as user preferences and business relevance can vary significantly between different regions, between urban and rural environments, and so on. Thus, all of our businesses of interest are in NYC, and our survey takers are either NYC residents or frequent visitors to NYC (this was used as a filtering question to be allowed to participate in the survey). We also limited the scope of our problem in terms of categories of interest. As we aim to construct a well rounded experience for a user visiting a source business, we limited ourselves to categories of businesses that people visit in person and that answer some recreational need. The list of categories we consider is shown in Table 1.

We asked the survey participants to tell us what information they find useful when they search for a specific business of some category c (in addition to information about the searched business). The answer options included:

1. Information about other nearby businesses from the same category

2. Information about nearby businesses from category c', for three candidate categories c' related to c

3. "None of the above"

The survey results are depicted in Figure 2a. The survey clearly showed that for any source category c, additional information about nearby businesses of some category $c' \neq c$ was found useful by a significant fraction of the survey takers. Furthermore, $c' \neq c$ was often the most useful category for additional information.

[1] For a small number of high-profile establishments (e.g. major tourist attractions like the Empire State Building), Google uses higher-precision data from its Knowledge graph to provide recommendations in its production system. Nethertheless, such establishments also appear like any other establishment in the RBG.

[2] http://www.google.com/insights/consumersurveys/home.

Indoor Lodging	Bars
Movie Theaters	Liquor Stores
Performing Arts Venues	Nightclubs
Sporting Goods Stores	Gyms
Beauty Salons	Restaurants and Cafes
Sports Complexes (e.g. Stadiums)	Tourist Attractions
Grocery Stores	Gift Shops

Table 1: **Categories of interest**. We only look at businesses in NYC that belong to one of the above categories as potential related business recommendations or as source businesses. Grocery Stores and Gift Shops are viewed only as potential recommendations, and we are not interested in them as source businesses.

We then examine how these user preferences match the current recommendations generated using Google‘s RBG. For every source category c, we examine the businesses x that belong to that category. We compare the fraction of x's top 20 RBG neighbors of category c to the fraction of neighbors in the top complementary category $c' \neq c$ of c according to the survey, and we aggregate these statistics over all businesses in category c. As portrayed by Figure 2b, in the vast majority of cases, the top RBG neighbors are businesses of the same category as the source business x, which does not match user preferences from the survey. There are only two cases where the complementary category c' is significantly popular among RBG neighbors of businesses of category c : $(c, c') = (Bar, Restaurant)$ and $(c, c') = (Nightclub, Bar)$. In both of these cases, the categories themselves are close to each other, and many businesses can be classified as both.

The above experiment demonstrates that Google's current related business recommendation system produces recommendations which are similar substitutes to the source business. However, users often prefer seeing nearby businesses of other, complementary, categories. This motivates our approach to provide more heterogeneous related business recommendation lists which include businesses complementary recommendations, in addition to substitutes.

4. ALGORITHMIC APPROACH

In this section we describe an algorithmic approach for generating a useful and diverse list of related businesses for a source business x. We assume that we are given the relevance scores between every pair of categories, and how relevant a business y is to x given a desired recommendation category. In the next section, we explain how we mine such signals from data, and for the purpose of the discussion in this section, we treat these two signals as given.

We first construct an optimization problem for generating a useful and diverse list of related business recommendations in the more general setting where each business can belong to any number of categories, and we explain the reasoning behind that construction. Since the problem under this setting is NP-hard, we relax the problem by associating each business to a single category, which results in an optimization problem that can be solved in polynomial time using a simple greedy algorithm.

Relevance Scores: We begin by explaining how the relevance between businesses is determined, and introducing notation that we will use in this section and throughout the rest of the paper.

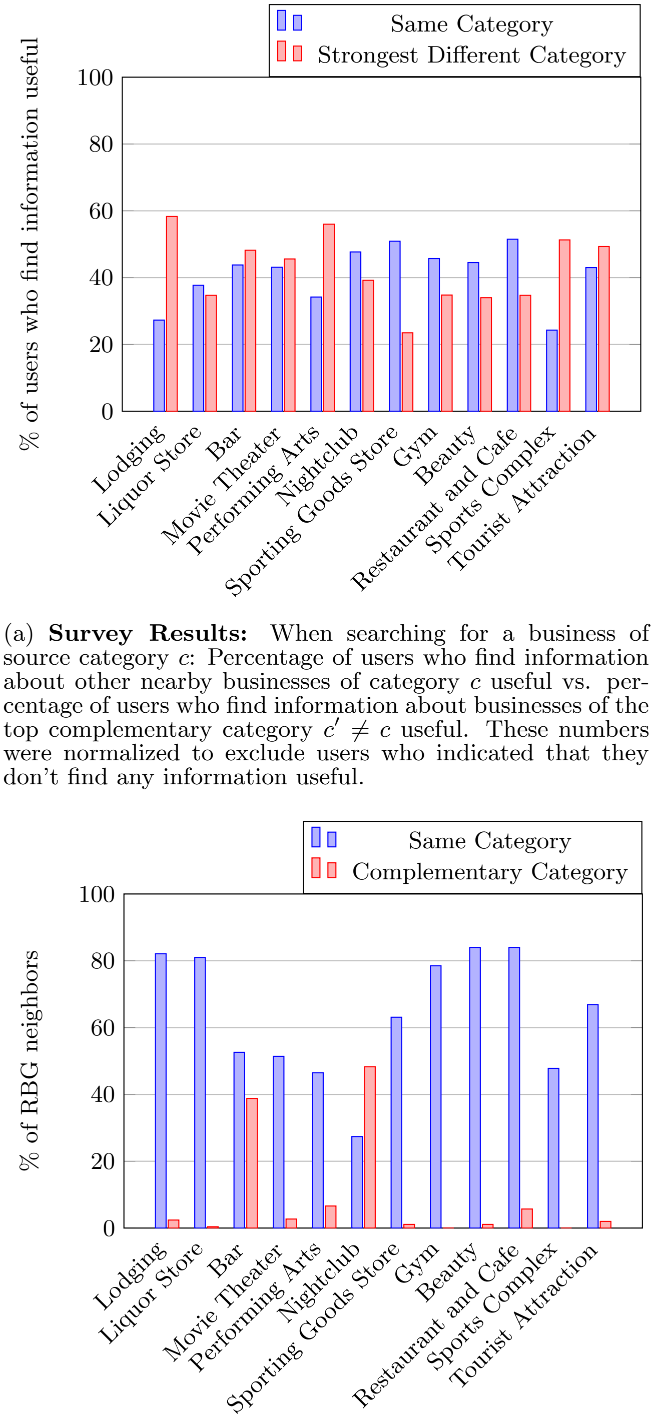

(a) **Survey Results:** When searching for a business of source category c: Percentage of users who find information about other nearby businesses of category c useful vs. percentage of users who find information about businesses of the top complementary category $c' \neq c$ useful. These numbers were normalized to exclude users who indicated that they don't find any information useful.

(b) **Top RGB neighbors:** For the top 20 neighbors of businesses of category c (aggregated): The fraction of neighbors of category c vs. the fraction of businesses of category $c' \neq c$ which was indicated most popular compilementary category for c by survey data.

Figure 2: **Same category vs. complementary category recommendations:** User preference compared to current Related Business Graph.

Let $r(y|x,c) \in [0,1]$ denote the relevance score of related business y to the source business x given that we are making a recommendation of category c. The relevance can be either relevance as a substitute if both x and y belong to category c, or as a complement otherwise.

Let Γ denote the set of all categories of interest, and let $R_{CC} : \Gamma \times \Gamma \rightarrow [0,1]$ be a function that gives the relevance between categories - that is, $R_{CC}(c_1, c_2) \in [0,1]$ denotes the relevance of category c_2 to category c_1 (the more relevant c_2 is to c_1, the higher the score is). Note that R_{CC} is not a symmetric function. Furthermore, R_{CC} is a function of categories only, and is agnostic of the source business. The semantics of R_{CC} are further discussed in the next section. As explained earlier, we treat r and R_{CC} as given in this section.

We define the relevance of a category c to a business x as the average relevance of c to the categories of x:

$$R_{BC}(x,c) = \frac{1}{\mid C(x) \mid} \cdot \sum_{c' \in C(x)} R_{CC}\left(c', c\right) \tag{1a}$$

where $C(x) \subseteq \Gamma$ is the set of categories to which x belongs. The relevance of a business y to x is then defined as:

$$R_{BB}\left(x,y\right) = \max_{c \in C(y)} \left\{R_{BC}\left(x,c\right) \cdot r\left(y|x,c\right)\right\} \tag{2a}$$

The relevance of y to x is determined through a single category in $C(y)$ - the one that maximizes the relevance to the source business - so a business's relevance is not penalized for belonging to multiple categories.

Optimization Problem - Multi-Category Setting: The output of our optimization problem is a list $S = (y_1, y_2, \ldots y_k)$ of k distinct businesses related to the source business x. One simple approach would be to display the k businesses that are most relevant to x under the above definition of business-to-business relevance. This approach would yield a list of related business recommendations where each recommendation y is highly relevant to x, but the list is a whole would be redundant. We would like to introduce diversity to the recommendation list by penalizing multiple recommendations from the same category.

The penalty for recommending multiple businesses from the same category should depend on the category, as redundancy may be worse for some categories than others. Also, to model diminishing returns in contributions made by recommendations of the same category, we would like the penalties for businesses in a category to be increasingly severe as we recommend more businesses from that category.

To satisfy these requirements, when constructing the list of related business recommendations $S = (y_1, y_2, \ldots y_k)$, we penalize multiple recommendations of the same category c by multiplying the business-to-business relevance terms with penalty terms of the form $d_c(rank_{c,S}(y))$. Here, y is a potential recommendation for category c, $d_c(n)$ is a category-c-specific *decay function* which is monotonically non-increasing, and $rank_{c,S}\left(y\right)$ is the number of businesses that occur before the index of y in S and belong to category c:

$$rank_{c,S}\left(y\right) = \mid \{j \in \{1,2,\ldots k\} \mid j < ind_S(y) \wedge c \in C(y_j)\} \mid \tag{3a}$$

$$ind_S\left(y\right) = \min(\{i \in \{1,2,\ldots k\} | y_i = y\} \cup \{k+1\})$$

(Note that the index of y in S is well-defined even if $y \notin S$, and that for such y, $rank_{c,S}(y)$ is just the number of businesses in S that belong to category c.) Adding the penalty terms gives us the final formulation of the optimization problem that yields a list of k related business recommendations that balances pairwise relevance to the source business x with overall list diversity:

$$Rec(x) = \underset{S=(y_1,y_2,\ldots y_k)}{\operatorname{argmax}} \sum_{i=1}^{k} \left[R_{BB}\left(x, y_i\right) \cdot \max_{c \in C(y_i)} \left\{d_c\left(rank_{c,S}\left(y_i\right)\right)\right\}\right] \tag{4a}$$

Note that the max operator over the possible penalty factors means that the penalty for adding a related business y to the list S is the least severe penalty out of all categories to which y belongs. The reasoning for this choice is that a potential related business should not be overly penalized because it belongs to multiple categories - for example, when recommending a bar as a related place for x, a bar-grill y that is classified as both a bar and restaurant should not be penalized just because other restaurants have already been recommended for x.

Optimization Problem - Single-Category Setting: Unfortunately, the optimization problem described in equation 4a, where every potential related business y can belong to a set of categories $C(y) \subseteq \Gamma$ of an arbitrary size, is NP-hard. We show NP-hardness of the problem using a reduction from the maximum k-subset intersection problem [16], and the proof is posted as supplementary material [1].

To enable a polynomial-time solution to the problem, we have decided to solve a relaxed version of the problem in which each business y is associated to a single category $c(y) \in \Gamma$. Belonging to multiple categories in Γ means that business y's (unique) immediate category is a subset of multiple categories in Γ under the given category ontology. For each such immediate category, we determined the most relevant superset category in Γ (see section 5 for further details).

After associating each business y to a single category $c(y)$, equations 1a to 4a takes the following form:

$$R_{BC}(x,c) = R_{CC}\left(c(x), c(y)\right) \tag{1b}$$

$$R_{BB}\left(x,y\right) = R_{BC}\left(x, c(y)\right) \cdot r\left(y|x, c(y)\right) \tag{2b}$$

$$rank_{c,S}\left(y\right) = \mid \{j \in \{1,2,\ldots k\} \mid j < ind_S(y) \wedge c = c(y)\} \mid \tag{3b}$$

$$Rec(x) = \underset{S=(y_1,y_2,\ldots y_k)}{\operatorname{argmax}} \sum_{i=1}^{k} \left[R_{BB}\left(x, y_i\right) \cdot d_{c(y_i)}\left(rank_{c(y_i),S}\left(y_i\right)\right)\right] \tag{4b}$$

A simple greedy algorithm, as described in Algorithm 1, finds an optimal solution to the relaxed optimization problem 4b in $O\left(\mid \Gamma \mid \cdot k\right)$ given a candidate pool of businesses sorted by their relevance score r (see online supplementary materials for proof of optimality and runtime).

Decay Functions To conclude this section, we discuss the decay functions d_c used for penalizing multiple recommendations of the same category. As mentioned before, these are monotonically non-increasing functions that model the diminishing returns of the contributions of same-category recommendations. The decay functions are category-specific,

Algorithm 1 Greedy algorithm for generating related business recommendation list under single category per. business relaxation

```
function Rec($x$)
    $S \leftarrow$ empty list
    while $| S |< k$ do
        $y^* \leftarrow \mathrm{argmax}_{y \notin S} R_{BB}\left(x, y\right) \cdot d_{c(y)}\left(rank_{c(y),S}\left(y\right)\right)$
        Append $y^*$ to $S$
    return $S$
```

$d_c^1\left(n\right) = \exp(-\alpha_c n)$
$d_c^2\left(n\right) = \exp(-\alpha_c n/2)$
$d_c^3\left(n\right) = \alpha_c^{-n}$
$d_c^4\left(n\right) \equiv 1$ (no decay)

Table 2: **Four decay functions**, from strongest to weakest, used in our experiments

and should provide a stronger decay, or steeper decrease, for categories in which multiple recommendations would cause highly diminished contributions. For this purpose, we introduce category-specific *decay factors:* $\alpha_c = R_{CC}\left(c, c\right)^{-1}$. The decay factor α_c is higher the less relevant a category c is to itself, which can be interpreted as how disadvantageous redundancy is among businesses recommendations of category c. The decay factors can then be used as a parameter for controlling the strength of the decay.

Table 2 lists the four families of decay functions we experimented with, from strongest to weakest. A strong decay function means that the algorithm would be fast to transition between categories, which can lead to a very diverse recommendation list. However, a strong decay function can also lead to transitions to categories with irrelevant businesses in them due to premature exhaustion of the most relevant categories. On the other hand, a weak decay function would make the algorithm stay within highly relevant categories, but may lead to an overly homogenous recommendation list. We will extensively compare the algorithm's empirical performance when using the different decay functions in section 6.

5. GENERATING RELEVANCE SCORES

In this section we explain how we constructed the signals R_{CC} and r which are used as inputs to the algorithmic framework presented in section 4. As mentioned before, $R_{CC}(c, c')$ measures relevance between of category c' to category c, and $r(y|x, c)$ measures the relevance of y to x as a recommendation of category c. These two signals were mined and composed from several data sources, including crowdsourcing, mobile maps queries, and the RBG edge weights. We begin this section with a description of the business category ontology used by the system, and then describe the construction of R_{CC} and r.

5.1 The Business Category Ontology

Google's business category ontology is based on a directed acyclic graph, G^{ONT} where vertices are categories and edges represent *"is-a"* relationships. An edge $(c_1 \rightarrow c_2)$ will be present if c_2 is a subset of c_1. For example, the edge (restaurant $\rightarrow$ chinese restaurant) is present in the graph. Every business is associated with a single immediate category which can be any vertex in G^{ONT}. However, for most businesses, the immediate category is quite specific.

We selected a set of categories Γ which is listed in Table 1. The categories in Γ are categories of recreational businesses, and are fairly general, so each $c \in \Gamma$ has many descendants in G^{ONT}. We would like to associate each immediate category c with all of its ancestors in Γ. However, the multi-category recommendation problem from section 4 is NP-hard, so for each immediate category c, we instead hand-pick the best ancestor of that category in Γ. As a result, each business is associated with a single category in Γ.

We were able to do this task manually because Γ is small, and because the number of immediate categories with multiple ancestors in Γ is also small (a few dozen). However, one could choose ancestors in Γ for each category using more scalable heuristics based on the topology of G^{ONT}, such as by choosing the closest ancestor in the graph.

5.2 Category-To-Category Relevance Scores

R_{CC} is a signal that measures the relevance between categories in Γ, such that $R_{CC}(c, c') \in [0, 1]$ is a score of how relevant category c' is to category c. $R_{CC}(c, c')$ measures how likely is a user visiting a source business of category c to visit another business of category c' (including the case of $c' = c$). Therefore, $R_{CC}(c, c')$ measures the extent to which c' *complements* c. We do not make any special additions to $R_{CC}(c, c)$ to account for increased relevance as a substitute, as the business-to-business scores $r(x|y, c)$, which are mostly based on the existing RBG, are already biased towards relevance as substitute.

R_{CC} is derived from three signals that give such category-to-category relevance scores:
(1) R^S is based on a crowdsourcing survey
(2) R^E is based on explicit mobile maps directions queries
(3) R^I is based on mobile maps search queries which can be treated as implicit directions queries

The signal R^S was computed using a crowdsourcing survey. The survey was conducted through the Google Consumer Surveys platform and included 500 participants who are residents or frequent visitors to NYC (similarly to the survey in section 3). The survey participants were asked to indicate which types of businesses they often *visit* in conjunction to a visit in a business of a specific category. For example, the survey participants were asked to answer the following: "When you visit a bar, which of the following types of establishments do you often visit in conjunction to that bar?". The answer options included:
1. Nearby businesses of the same category (e.g. nearby bars)
2. Nearby businesses of three candidate related categories (e.g. nearby restaurants or nightclubs)
3. I don't visit businesses of that category
4. None of the above - i.e. I visit businesses of that category, but I don't typically visit other businesses in conjunction

$R^S(c, c')$ is then the fraction of survey takers who indicated that they visit a business of category c' when visiting a business of category c. R^S was normalized to exclude survey takers who chose answer 3, but to include survey takers who chose answer 4. We take answer 4 to indicate that showing complementary recommendations for source category c is less beneficial, so we diminish the magnitude of $R^S(c, c')$ for all potential related categories c' (including $c' = c$).

The signals R^E and R^I are generated using data from Google's mobile maps of both explicit (R^E) and implicit

(R^I) directions queries. An explicit directions query is a query where a user asks for directions from business x to business y on the mobile maps application. Unlike a directions query, in a mobile maps search query, a user searches only the location of the business on the mobile maps application. An implicit directions query from x to y is a pair of such maps search queries for x and then for y issued by the same user within a short time window. Both for data reliability and robustness, and to avoid any potential privacy concerns, when computing these signals, we only look at aggregated counts for pairs of businesses (x, y) that occur sufficiently often in the dataset. Before aggregating by source and destination categories, our dataset included 1 million explicit directions queries and 10 million implicit directions queries. Direction queries are a strong indicator of complement relations between businesses - i.e. if there are many directions queries between x and y, then this is an indicator that y complements x.

Directions queries are highly biased towards queries by tourists and other visitors, so hotels, restaurants, and tourist attractions are by far the most common categories associated with almost any category. Therefore, simply mining association rules using commonly occurring pairs of source and destination categories would not work, so we must consider another metric. We use a notion similar to tf-idf for extracting relevance from the directions query dataset. If E is the multi-set of category-annotated explicit queries $q = (c_{src}, c_{dst})$, then we define R^E as follows:

$$R^E(c, c') = \frac{\mid \{q \in E | q.c_{src} = c \wedge q.c_{dst} = c'\} \mid}{\mid \{q \in E | q.c_{src} = c\} \mid +1} \times \log\left(\frac{\mid E \mid}{\mid \{q \in E | q.c_{dst} = c'\} \mid +1}\right)$$

We define R^I in a similar way using the multi-set of implicit queries.

We combine R^S, R^E, and R^I into a single signal R_{CC} by taking a weighted average. R^S is sparse, as only 48 out of the 168 possible directed pairs of categories were included in the survey. Thus, R^S is only counted towards R_{CC} for pairs of categories included in the survey $(c, c') \in Srv$. R^S is based on answers by people who live in NYC, or visit there frequently, whereas the directions queries are biased towards tourists and other non-locals. To account for both crowds, we take the average of R^S and a signal composed of R^E, and R^I. Since explicit directions queries are higher precision and a clearer signal of user intent, we weight each explicit query twice as much as an implicit query, and since we have 10 times as many implicit as explicit queries, we average R^E and R^I in a ratio of 1:5. This gives us the final formula for R_{CC}:

$$R_{CC}\left(c, c'\right) = \begin{cases} \frac{1}{2} \cdot R^S\left(c, c'\right) + \frac{5}{12} \cdot R^I\left(c, c'\right) + \frac{1}{12} \cdot R^E\left(c, c'\right) & (c, c') \in Srv \\ \frac{5}{6} \cdot R^I\left(c, c'\right) + \frac{1}{6} \cdot R^E\left(c, c'\right) & (c, c') \notin Srv \end{cases} \tag{5}$$

In the approach described in section 4, the signal R_{CC} is used for both establishing business-to-business relevance through category adjustment, as well as to control for the strength of the decay functions. Typically, our approach produces a more diverse related business recommendation list than the current Google RBG-based system does. However, in cases $R_{CC}(c, \widetilde{c})$ is very high, and $R_{CC}(c, c')$ is low for all $c' \neq \widetilde{c}$, our system might in fact un-diversify the recommendation list by making recommendations only, or mostly, from category $\widetilde{c}$. As we will explain in section 6, this case did come up in our evaluations.

5.3 Business-To-Business Relevance Scores

$r(y|x, c)$ is a signal that measures the relevance of a potential related business y to the source business x given that we want to make a recommendation of category c. To compute $r(y|x, c)$, we combine three signals that give business-to-business relevance scores:

1. $r^{RBG}(y|x)$ is the normalized weight of the edge (x, y) in Google RBG: $r^{RBG}(y|x) = \frac{w(x,y)}{\sum_{y'} w(x,y')}$
2. $r^E(y|x)$ is the fraction of queries with destination y among explicit mobile maps directions queries from x
3. $r^I(y|x)$ is the fraction of queries with destination y among implicit directions mobile maps queries from x

For each of the above three $r(y|x)$ signals, we define $r(y|x, c)$ to be $r(y|x)$ if y is of category c, and zero otherwise.

r^{RBG} has far greater coverage of possible business pairs than r^I does. On the other hand, directions queries give more obvious and direct evidence of business-to-business relevance and complementarity. So r^{RBG} is a higher recall and lower precision signal than r^I, which is in turn a higher recall and lower precision signal than r^E. We combine a higher recall and lower precision f_1 with a lower recall and higher precision signal f_2 by using f_2 as an amplifier for the f_1 scores: $combine(f_1, f_2) = f_1 \times (1 + f_2)$. r^{RBG} is used as the basis for r, and we successively combine each signal with the next high recall signal, so we get the final formulation for $r(y|x, c)$:

$$\begin{aligned} r\left(y|x, c\right) &= combine(r^{RBG}, combine(r^I, r^E)) \qquad (6) \\ &= r^{RBG}\left(y|x, c\right) + r^{RBG}\left(y|x, c\right) \times r^I\left(y|x, c\right) \\ &+ r^{RBG}\left(y|x, c\right) \times r^I\left(y|x, c\right) \times r^E\left(y|x, c\right) \end{aligned}$$

Locality is implicitly a factor in $r(y|x, c)$. The RBG takes into account locality, so high geographic proximity between x and y will increase the weight of the edge (x, y), and thus increase $r\left(y|x, c\right)$.

6. EVALUATION

We evaluate several variants of our system using human raters, and compare them against Google current RBG-based system.

The Evaluation Procedure: The raters are professionals who undergo training by Google for ranking results, and are experienced in side-by-side comparison of alternative business result sets. If a rater is not familiar enough with any displayed business, she is instructed to research that business. We gave the raters items that showed a source business x and two lists of 5 related businesses each - one produced by selecting the top neighbors of x in the RBG, and the other produced by our system. The instructions for the item asked the raters to assume that a user had searched and was considering visiting x. The side that our list appeared on was randomized.

For each item, we gave the raters two tasks. The first was to rate how useful each related business recommendation y was for the user. For each recommendation y, the raters had to choose one of three options:
(1) ‘Useful as a Complement’ - y is a useful suggestion for a complementary business that the user could visit in conjunction to x
(2) ‘Useful as a Substitute’ - y is a useful suggestion for a substitute that the user could visit instead of x
(3) ‘Not Useful’
The second task was to compare the two lists for overall usefulness for the user. The raters had to choose whether the list generated by our system was Much Better / Better / Slightly Better / About the Same / Slightly Worse / Worse / Much Worse than the list generated by the current RBG-based system. Furthermore, they had to give a one sentence explanation for their rating.

We selected 500 items by choosing the 500 businesses with the largest number of explicit directions queries with that business as an origin. Since the explicit directions queries is our sparsest data source, this approach allowed us to utilize as many of our available data sources as possible. We had each of these items rated by three independent raters, for a total of 1500 comparison ratings per evaluation procedure.

We repeated this evaluation procedure for a variant of our system using each of the decay functions listed in Table 3. The choice of decay function had a significant impact on the recommendation lists, and we found that the stronger the decay was, the further our system's recommendation lists were from the RBG-based lists. Table 3 shows the average set overlap between our system's recommendation list and the RBG-based list for the same source business for each decay function.

Overall List Ratings: We start by analyzing the results of the raters' second task, as this task's ratings are the indicator of our system's relative performance compared to that of the current Google system. We explore the ratings for the lists generated by our system for each of the decay options. These ratings compare the lists against the equivalent lists generated by the current Google RBG-based system and give the gain or loss of our system compared to the current system. The distribution of the results is described in Figure 3. We assign a numeric score to each rating option — from (-3) for ‘Much Worse’ to (+3) for ‘Much Better’, and display the statistics of the improvement compared to the current RBG-based system in Table 4.

For all of the decay options, the lists generated by our system received a positive average rating (with a confidence level of at least 95%). This means that our system achieves a statistically significant improvement in recommendation list ratings, and that on average, our system produced a more useful recommendation list than the original RBG-based system.

The choice of decay function has a significant impact on the rating distribution. A strong decay function may lead to premature exhaustion of relevant categories and switching to irrelevant categories. Such a scenario usually leads to strong losses, which are more common when using strong exponential decay. On the other hand, using no decay would commonly lead to an homogenous recommendations list, missing some of the opportunities for high gain that arise from diversifying the list. An intermediate decay function, $d_c^3(n) = \alpha_c^{-n}$, achieved the best improvement in the generated lists.

Decay Function	Average Overlap with RBG list
$d_c^1(n) = \exp(-\alpha_c n)$	37%
$d_c^2(n) = \exp(-\alpha_c n/2)$	46.5%
$d_c^3(n) = \alpha_c^{-n}$	51.1%
$d_c^4(n) \equiv 1$ (no decay)	69%

Table 3: **Decay functions and average overlap:** The four decay functions we explored, from strongest to weakest, and the average set overlap between a recommendation list generated by our system using that decay function and the original recommendation list for that same source business as generated from the RBG.

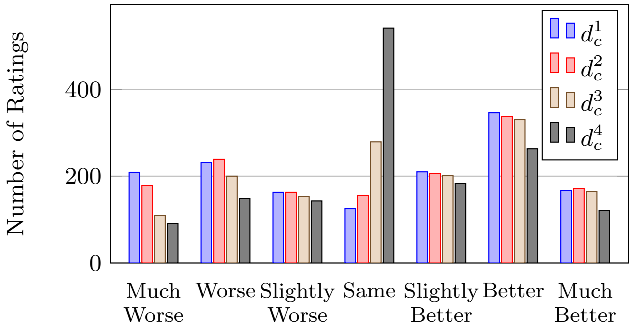

Figure 3: **Comparison to RBG - overall rating distribution:** For each decay function, the distribution of the comparison scores between lists generated by our system using that decay function and lists generated from the RBG.

Individual Business Ratings: Next, we analyze the results of the raters' first task, where they had to classify each recommendation as a useful complement, useful substitute, or neither with respect to the source business. The results are listed in Table 5. As expected, the recommendations generated by the original RBG-based system have a strong majority of useful substitutes and a small fraction of useful complements, but also a small fraction of not useful outliers. The stronger the decay is, the farther the rating distribution is from that of the RBG-based system, and the stronger the decay is the lower is the fraction of useful substitutes among the recommended businesses, and the higher are the fractions of useful complements on one hand and the fraction of not useful suggestions on the other. Note that even when no decay is used and there is no incentive to switch between categories, the fraction of useful complements is higher in our system compared to the RBG-based system, and the fraction of the substitutes is lower. This is mostly due to cases where the most relevant category according to R_{CC} is a category different than the category of the source business, as the RBG-based system tends to make recommendations similar to the searched business.

Per Category List Ratings: The items that we have analyzed are for the 500 most popular source businesses in the explicit directions queries dataset. This dataset has a strong bias towards tourists and other non-locals, so the set of 500 businesses includes many lodging options (258), restaurants and cafes (94), and tourist attractions (72). We now analyze the performance of our system by *source category* (i.e. the

Decay	Average Rating	95% Confidence Interval	positive ratings	negative ratings
d_c^1	0.10	0.0-0.21	49.8%	41.6%
d_c^2	0.15	0.04-0.25	49.2%	40.0%
d_c^3	0.33	0.23-0.42	48.4%	32.2%
d_c^4	0.24	0.16-0.32	38.0%	25.7%

Table 4: **Comparison to RBG - rating statistics:** Comparison of the lists generated by our system to the original RBG-based lists. The ratings are in the scale of (-3) to (+3) with a positive rating indicating improvement compared to the original system, and a negative rating indicating a loss. For all decay options, the average rating is positive with a confidence level of 95%

Decay	Complement	Substitute	Not Useful
d_c^1	43.6%	32.3%	24.1%
d_c^2	36.6%	42.1%	21.3%
d_c^3	35.8%	45.4%	18.8%
d_c^4	21.5%	64.4%	14.1%
RBG	14.4%	71.6%	14%

Table 5: **Distribution of individual recommendation ratings:** For each variant of our system, and for the original RBG-based system. Every recommendation was classified by the raters as either 'Useful Complement', 'Useful Substitute', or 'Not Useful'.

category of the searched source business) to see when is our approach the most applicable.

Since indoor lodging, restaurants and cafes, and tourist attractions are the only categories represented in significant numbers in the previously evaluated business set, we ran separate evaluations for the rest of the source categories listed in Table 1. For each of the remaining source categories, we selected the most popular 50 source businesses of that category in the explicit directions dataset. The results of comparing the list generated by our system using d_c^3 as a decay function against the equivalent lists generated by the current RBG-based system appear in Table 6.

Our system achieved a significant improvement in recommendation list usefulness for the beauty, indoor lodging, liquor store, and sporting goods store. Our system achieved an improvement, but failed to do so with a 95% confidence level, for the bar, nightclub, performing arts venue, movie theater, gym, and sports complex categories, and suffered small losses on average for the restaurant and tourist attraction categories. To reason about the cases in which our system achieved improvement, we also address the ratings of individual recommended businesses (i.e 'useful complementary' / 'useful substitute' / 'not useful') which are not listed here for the by-category evaluations for space considerations.

In the case of the beauty category, our system actually recommends more useful substitutes than the RBG-based system does. The high fraction of useful substitutes in the suggestions is due to the high relevance of beauty to itself, i.e. high $R_{CC}\left(beauty, beauty\right)$, and the low relevance of other categories to beauty. The category adjustment encourages the system to start with other beauty recommendations, and not to quickly switch to other categories due to a low decay factor. Furthermore, as businesses in the beauty category are often geographically sparse, the RBG-based system, which considers distance more directly than our system does, would often recommend similar but not directly related businesses such as clothing stores that are close geographically. Raters often do not find such loosely related recommendations useful, and harshly penalize the RBG-based list. A similar scenario explains the gain our system achieved for the liquor store, sporting goods store, and to a lesser extent, the bar and nightclub categories. All of these five categories have very high self-relevance scores $R_{CC}\left(c, c\right) > 0.75$.

Category	Average Rating	95% C.I.	positive ratings	negative ratings
Beauty	0.71	0.41-0.99	53.7%	23.1%
Indoor Lodging	0.52	0.38-0.66	59.4%	34.4%
Liquor Stores	0.37	0.15-0.58	29.9%	13.6%
Sport Stores	0.33	0.1-0.54	33.3%	19.7%
Sports Complexes	0.26	(-0.04)-0.54	45.1%	34.7%
Bars	0.18	(-0.07)-0.42	39.4%	28.6%
Nightclubs	0.15	(-0.09)-0.37	30.6%	23.1%
Performing Arts	0.10	(-0.13)-0.34	33.3%	29.9%
Movie Theaters	0.07	(-0.16)-0.29	29.9%	25.9%
Gyms	0.06	(-0.12)-0.24	18.8%	20.1%
Restaurants	0.00	(-0.16)-0.15	29.9%	27.0%
Tourist Attractions	-0.01	(-0.25)-0.21	38.3%	37.8%

Table 6: **Overall list ratings organized by source category** - comparing the lists generated by our system, when using d_c^3, to the original RBG-based lists. The categories are sorted from highest to lowest improvement, as indicated by the average rating.

In the case of the indoor lodging category, our system suggested significantly more useful complementary businesses (45.1%) than the the RBG-based system did (6.8%). For example, the related businesses recommendations the RBG-based system provides for the source business "Hotel 373 Fifth Avenue" - a boutique in Midtown East, Manhattan are as follows: $(y_1, y_2, y_3, y_4, y_5) =$ (Hotel Giraffe, Courtyard NY Manhattan-Fifth Avenue, New York Marriott Marquis, Courtyard Times Square South, 70 Park Avenue: A Kimpton Hotel) - all are Midtown hotels similar to the source business. On the other hand, our system makes the following recommendation list: (Hotel Giraffe, American Girl Place Cafe, Courtyard NY Manhattan-Fifth Avenue, Bryant Park, Tenpenny) - i.e. the list is a mix of two similar hotels, two restaurants, and a famous park - all in Midtown East. The raters favored the diversity, even at the expense of a higher fraction of not useful recommendations. A similar scenario explains the improvements that our system achieved in the cases of the sports complex and performing arts venue categories. For all three categories, survey takers in the survey described in section 3 indeed indicated a strong preference for information about complementary businesses when searching for a source business belonging to one of these categories.

In the cases of movie Theater and gym in which our system achieved a very small improvement, the distributions of individual recommendation ratings is very similar for both our system and the RBG-based system. In the cases of restaurant and cafe and tourist attraction categories, our system recommended less useful substitutes, more useful complements, and slightly more not useful recommendations, and suffered a small loss compared to the RBG-based system. Some of the losses are due to locality being less explicitly taken into account in establishing business-to-business relevance compared to the RBG, and some recommendations lose relevance drastically if not immediately close to the source business (e.g. recommending a gift shop not immediately close to the searched tourist attraction). A future improvement to our system could be including distance-based regularization that could more directly account for locality.

7. CONCLUSION AND FUTURE WORK

We studied the problem of improving the usefulness of the list of related business recommendations a search engine provides in response to a specific business search. Our approach was to build a relevant and diverse mixture of substitute and complement recommendations. In our problem settings, an ontology of business categories exists. We mined multiple data sources to establish business-to-business relevance scores that take that ontology into account. We introduced an algorithm that constructs a list of related business recommendations for a given source business. The algorithm balances pairwise business-to-business relevance scores with diversification through penalizing multiple recommendations from the same category. We explored several penalization options through different families of decay functions.

Compared to the current Google system, our system typically generates recommendation lists with more complementary business recommendations and fewer substitutes similar to the source business. The evaluation results show a significant improvement in recommendation list usefulness for the lists generated by our system compared to the current Google system, which demonstrates the added value in including local complementary businesses recommendations. However, the extent to which a source business would benefit from complementary vs. substitute recommendations varies greatly on the source business and its category. In some cases, our system actually gained user satisfaction through offering more substitutes similar to the source business.

Furthermore, in this work we treat two businesses as substitutes if they belong to the same category, and as complements if they are otherwise relevant to one another. However, this is just one interpretation for substitute vs. complementary. Another interpretation could be clustering categories and determining substitute relation iff two businesses are in categories belonging to the same cluster. Other possible interpretations of business to business substitute vs. complement relations can be completely unrelated to the ontology and rely on entirely different features or data sources (e.g. online reviews or social network based connections).

Lastly, our work did not take into account user-specific features, and addressed providing recommendations to an anonymous random search engine user. However, user settings could and should be taken into account, even for providing non-personalized recommendations. In particular, the system could use the user's location and whether they are using a mobile or non-mobile device to provide the best recommendations for a well-rounded and satisfying experience wherever they may go.

The factors that make two business substitute or complement each other, the different interpretations for such relations, and the circumstances in which substitutes vs. complements should be recommended, are all possible lines for future work on the topic.

Acknowledgments. We thank the local search quality team at Google New York, and to Mayur Thakur in particular, for their support and feedback. We also thank Jure Leskovec and the SNAP group at Stanford, and to Austin Benson and Julian McAuley in particular, for their useful feedback. Additionally, we thank Tilman Achberger for his help with the evaluation process. Yonathan Perez is supported by P. Michael Farmwald Stanford Graduate Fellowship.

8. REFERENCES

[1] Supplementary Materials. https://www.dropbox.com/sh/wau25zqwusnv8mc/AADJta6Ybc6RzyJiJRRCCk_ea?dl=0.

[2] R. Agrawal, S. Gollapudi, A. Halverson, and S. Ieong. Diversifying search results. In *WSDM*, 2009.

[3] J. Carbonell and J. Goldstein. The use of mmr, diversity-based reranking for reordering documents and producing summaries. In *SIGIR*, 1998.

[4] C. Clarke, M. Kolla, G. Cormack, O. Vechtomova, A. Ashkan, S. Buttcher, and I. MacKinnon. Novelty and diversity in information retrieval evaluation. In *SIGIR*, 2008.

[5] V. Dang and B. Croft. Diversity by proportionality: An election-based approach to search result diversification. In *SIGIR*, 2012.

[6] M. Drosou and E. Pitoura. Search result diversification. *SIGMOD Record*, 39(1):41–47, 2010.

[7] G.Linden, B.Smith, and J.York. Amazon.com recommendations: Item-to-item collaborative filtering. In *IEEE Internet Computing*, 2003.

[8] S. Gollapudi and A. Sharma. An axiomatic approach for result diversification. In *WWW*, 2009.

[9] J. Levandoski, M. Sarwat, A. Eldawy, and M. Mokbel. Lars: A location-aware recommender system. In *ICDE*, 2012.

[10] J. McAuley, R. Pandey, and J. Leskovec. Inferring Networks of Substitutable and Complementary Products. In *KDD*, 2015.

[11] D. Rafiei, K. Bharat, and A. Shukla. Diversifying web search results. In *WWW*, 2010.

[12] R. Santos, C. Macdonald, and I. Ounis. Exploiting query reformulations for web search result diversification. In *WWW*, 2010.

[13] S. Vargas, L. Baltrunas, A. Karatzoglou, and P. Castells. Coverage, redundancy and size-awareness in genre diversity for recommender systems. In *RecSys*, 2014.

[14] E. Vee, U. Srivastava, J. Shanmugasundaram, P. Bhat, and S. Amer-Yahia. Efficient computation of diverse query results. In *ICDE*, 2008.

[15] P. Venetis, H. Gonzalez, C. Jensen, and A. Halevy. Hyper-local, directions-based ranking of places. In *VLDB*, 2011.

[16] E. Xavier. A note on a maximum k-subset intersection problem. *Information Processing Letters*, 112(48):358–373, September 2011.

[17] M. Ye, P. Yin, and W. Lee. Location recommendation for location-based social networks. In *GIS*, 2010.

[18] C. Yu, L. Lakshmanan, and S. Amer-Yahia. It takes variety to make a world: Diversification in recommender systems. In *EDBT*, 2009.

[19] C. Zhai, W. Cohen, and J. Lafferty. Beyond independent relevance: Methods and evaluation metrics for subtopic retrieval. In *SIGIR*, 2003.

[20] M. Zhang and N. Hurley. Avoiding monotony: Improving the diversity of recommendation lists. In *RecSys*, 2008.

[21] J. Zheng, X. Wu, J. Niu, and A. Bolivar. Substitutes and complements: Another step forward in recommendations. In *EC*, 2009.

[22] C. Ziegler, S. McNee, J. Konstan, and G. Lausen. Improving recommendation lists through topic diversification. In *WWW*, 2005.