Context Ontologies for Recommending from the Social Web

Eoin Hurrell
CLARITY: Centre for Sensor Web Technologies
and School of Computing
Dublin City University, Glasnevin, Dublin 9
ehurrell@computing.dcu.ie

Alan F. Smeaton
CLARITY Centre for Sensor Web Technologies
and School of Computing
Dublin City University, Glasnevin, Dublin 9
alan.smeaton@dcu.ie

ABSTRACT

Investigations into combining context and recommendation has resulted in much fruitful research which has improved recommender systems. Such contextual information has come in many forms and been used in different ways, successfully offering better in-situ suggestions. Factors such as location, time of recommendation, etc. have proven themselves as useful contributors to exploiting context. One issue, however, is the importance placed on each aspect of context, especially as new forms of recommendation are developed. Context is traditionally incorporated into recommenders at design-time, as a filter or as an integral part of how users are modelled, but the importance placed on each aspect is not often examined. Social recommenders and systems that draw on the wealth of data present in social networks frequently have access to far more contextual factors than traditional recommenders, making user relationships to these factors all the more important. The main contribution of this paper is to provide an examination of contextual priorities from the social web, which prove useful to recommender research in the area. This ontological examination of context shows that users have different priorities when it comes to context with a large variation in the suitability of each contextual factor in predicting good recommendations. In addition, this paper presents and discusses an approach to individually tailoring context ontologies (allowing for dynamically generated context sets), evaluating contextual factors in recommending from the social web.

Categories and Subject Descriptors

H.3.3 [**Information Search and Retrieval**]: Information Filtering; H.1.2 [**User/Machine Systems**]: Human Factors

Keywords

context, social web, user behaviour, recommendation

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
ACMWSDM '12 Rome, Italy
Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...$15.00.

1. INTRODUCTION

Contextual recommendation systems are designed to meet the needs of people by capturing and leveraging the factors that currently influence them. Thus it is important to try and understand which context factors or aspects contribute to decision-making and whether these are generally shared across different users. With the widespread use of social networks, and the growing use of recommender systems leveraging the data they make available, it is important to study context within the domain of such social networks.

We wish to explore how useful people find different forms of context, and whether an ontological model of their degree of interest in such factors can be generated. Until recently it has been assumed that in contextual recommendation, all forms of context available should be used for any task. Recent work has shown that some contextual information is not relevant for some tasks [3], and here we similarly examine the idea that a contextual feature is agreed on by users to be important for recommending within the domain of the social web. We investigate whether the redundancy of some contexts exists in the social web and whether it can be detected dynamically. If so do individual users have preferences for distinct context sets within a system? In essence we looked at what factors had the greatest impact on whether a person followed someone on a social network, indicating priorities people place on contexts when they are making decisions.

Any contextual features that could account for situations where otherwise good suggestions could be rejected should be of interest. The fact that memory-based recommender systems focus on forming groups of users from what is known about them, essentially stereotyping people, creates challenges and the more we know in the form of contextual data, the harder it is to decide how to form groups. In essence there is a risk of creating a "contextual long tail" by stating that users only share an interest if that interest is rated highly in the exact same context. Contextual data might be important by design for the given task, or different information might be important to different people, hence our undertaking to study this area.

In this paper we detail a study in which we examined the contextual factors surrounding the relationships between a social network user and the people that they follow. We do this as a method of examining the contexts that impact the task of friend recommendation within as social network. This demonstrates an approach to validate the impact of context across users and shows the ability of different contexts to predict relationships.

2. RELATED WORK

Context is broadly understand to mean the intrinsic and extrinsic information surrounding something that contributes to our understanding of it. The concept of context is integral to human communication; it forms the basis of many language constructs that allow people to understand who is being referred to by pronouns for example. When dealing with computing systems, context is simply any information that tells us more about the user or or the processes or tasks they are engaged in, including sensed information such as time or location. The definition of context has grown to encapsulate complex semantic interplays between an item and its environment. People have a cultural understanding of context related to how they use language [12] and a person's context can be said to be anything that affects that person's decisions, as shown in [30]. It is also worth noting that social concepts of public and private, such as those related to the sharing seen in recommender systems, have always been intimately tied to representations of context such as location (discussed, for example, in [26]). Locations define the sort of interactions that are appropriate within them by virtue of how public or private they are, from the privacy of one's home to the public space of a large shopping centre. Savolainen [23] shows context is a key factor in "everyday life information seeking". For all these reasons it is desirable to examine the ways in which recommender systems can understand the importance of different contexts and context combinations, for users.

Context in computing has been the subject of much research as shown in the review by Dey [10], and integrating context into recommendation has proved useful. Work has been done by Ingwersen and Järvelin [15] to define the variety of contexts that exist, independent from their means of collection. This work was refined by Ingwersen [14] who shows a range of contextual factors that cover things that can be directly sensed, such as location (using GPS) or time, and 'offered information', such as a person's name. Sensed information, inferred from physical sensors and used in such as applications as location-based recommendation, provides a means of detecting some of these contexts.

A set of seven contextual layers has been defined in [15] as follows:

Intra-object context This context relates to the relationship an object has with itself. It can involve metadata and the connections between item attributes, or the quantifiable structure of the item, particularly of textual content. In social networks the profile information users see when evaluating others falls into this category.

Inter-object context This encompasses all the factors involved with relations between items, assigned index terms or external metadata that relates to the item. Playlists are a good example of this, as they connect items in a context that they would not have on their own.

Session context Session context is the context gathered from a single usage or session, a person's usage patterns in the recommender, which involves real user tests or interaction simulation. This is the most common context used by recommender systems, current location, time and so forth and

Individual contexts This relates to the social, conceptual, emotional or systematic contexts specific to the user. Their impact can be seen in rating behaviour and usage. Social networks, with tasks like friend recommendation, blur the line between this and "Intra-object" context, as the objects being recommended are users, giving us the opportunity to study how people evaluate context as we do here.

Collective contexts This relates to the social, conceptual, emotional or systematic contexts the user inherited from the collective, be it through membership of a community or though being grouped with like-mined users. Though recommendation frequently involves grouping users, contextual recommender research has not focused on varying context usage based on these groups. In our work we will touch on this by examining whether there is a communal common interest in context within social networks..

Techno-economic and societal contexts These, somewhat more global contextual factors affect all previous contexts, but in ways that can be difficult to detect.

Historical contexts Historical context refers to the collection of previous events that could influence a person's decision-making.

Currently contextual recommendation makes good use of a number of common features, prompting some research discussion for example in work by Schmidt *et al.*[24]. Contextual recommendation work has identified three ways to integrate context into recommendation ([1]), filtering items before (pre-filtering) or after (post-filtering) recommendation occurs and altering the representation of user-item ratings to be user-item-context ratings. Each has been considered to have advantages and disadvantages, while comparisons as described in [21] have shown that neither pre- nor post- is significantly better, resulting in designs for context in recommenders that are usually decided at build time, with little study of how contexts are actually used by individuals for the application.

The development of social networking sites such as Twitter[1] and Facebook[2] has lead to an explosion of new data being shared by people, in effect giving others more information with which to form decisions about them. These new contextual factors can range from where they post updates, what device they use, to whether they use a pseudonym and how they style their profile page. This is in stark contrast to the relative sparsity of contextual data that lead to the common use of location and time within contextual recommendation.

Derrida once famously said "There is nothing but the context" [9], highlighting the importance of understanding surrounding factors in understanding the person. Accounting for context in recommendation is hugely desirable, as we have shown, and research suggests that it improves accuracy, already pointing out that context is of value in harnessing the explosion of additional information brought about by the realtime social web [19].

Recent work by Google reports that 70% of smartphone owners use their device while shopping, and the majority

[1]http://twitter.com
[2]http://facebook.com

of shoppers use online resources for research and purchase in their local store [13]. Mobile applications have now been developed that prove the viability of item suggestion in a mobile context [6], and of using location to inform suggestions [29, 5, 22]. These factors point to a future of computing in a retail context that will benefit from the personalisation opportunity and interaction offered by a recommender that is contextually relevant and aware. Research by Schmidt *et al.* [24] warns against the focus on location as a quick and easy contextual factor while missing out on the multitude of other contexts, both sensed and surveyed, which are possible. Interestingly, most contextual recommendation work treats contexts as continuous variables, while work by Anand *et al.*[2] shows that discreet "finite states" also work, but have not been widely studied. Here we will investigate which method(s) users prefer when expressing context.

It is far from simple to use context, as Dhar *et al.* showed that even time pressure for example has a huge effect on other contextual features and how they are perceived [11]. As previously mentioned, recent research [28] has defined three major methods for incorporating context into recommendation algorithms. These three methods are pre-filtering, post-filtering and altering the user model. The drawbacks of these methods in traditional recommendation is that none provide a method to determine which contextual factors are of primary importance dynamically, which is what we study here. Since CF recommenders work by forming groups based on user information, any new information has the potential to further subdivide groups, and since recommendation quality is directly related to the size of these groups, context must be intelligently managed. In essence there is a risk of creating a "contextual long tail" by stating that users only share an interest if that interest is rated highly in the exact same context, which has gone without study.

Ingwersen and Järvelin [15] argue for a breaking down of the division between quantitatively-oriented IR and qualitatively-oriented IS in the study of context due to its complexity and dependence on user sentiment. We explore in this work both qualitatively and quantitatively, the attitudes users have toward context and the use of multiple context factors in recommendation.

The place of features such as sensed context (then considered as part of a measure of performance) has been debated since before sensing technology became as sophisticated as it currently is [18]. Here we show it is possible to measure the performance of represented contexts such as place, time and the online identity as features for each user of a system.

It is well known that choice is affected by context, investigated by Yoon and Sominson [30] and by Dhar *et al.* [11], which could be for a number of reasons, perhaps relating to perceived inconvenience, tying in with the work reported in [8]. As has been mentioned earlier, only some contextual features are relevant for any given decision within recommendation [3], and work by Madani and DeCoste highlights that not all context impacts recommendation [17]. Here we turn our attention to user-level contextual feature selection, finding that each user is indeed different in the features they consider. In the past designing for context has been styled as scenario oriented recommendation, in that recommenders are then only useful in the envisioned scenarios [25].

To further delve into the relationship between context and recommendation, recommender systems built to be "context-aware" such as discussed in [1] would further benefit from being "user-aware" in the choice of that context, as we investigate here. Machine learning is not new in recommendation [4], but here we apply it in a novel way to the domain of social networks. Previously, contextual recommendation has used a single Support Vector Machine (SVM) to model context over all users [20]. Here we train an SVM for each user to examine how each user benefits from each feature. We do this for much the same reason as Noulas *et al.* conducted their research into modelling context using random walks [19], addressing the problem of an abundance of contextual data available to improve recommendation becoming available from a variety of sources. This work can be seen as an extension of work by Koren [16] into latent factors in recommendation, but applied to the new area of contextual factors on the social web. What results is a quantitative analysis of the priority individual users place on different forms of context.

3. METHODOLOGY

Our experiment is designed to highlight the contexts people are interested in when following a user on Twitter[3]. Twitter is a social network micro-blogging site that allows a user, under a screen name, to compose 140-character messages for people who are following them, to read. Users have followers and friends who they themselves follow to see updates (called tweets) from. Other features like marking a tweet as "favourited", putting users in lists and "retweeting" (forwarding a message from someone else to all of your own followers) also exist. Many of these user-generated microblog streams are publicly available. Users within this system represent a dense collection of contextual factors, making it desirable to understand how people evaluate each of these contexts.

We collected a dataset of tweets from publicly-accessible twitter users, using the "firehose" Twitter API. We gathered 251,807 tweets from 7,390 unique Twitter users within the Dublin area. We restricted our collection of tweets to one area in order to control for a number of factors. The most obvious was timezone as we examined, among other contexts, the times people tweeted. There was also a concern that contextual evaluation might vary across cultural or social divides, hindering the prospect of finding common contextual features.

Twitter provides a wealth of data with each tweet. We took 61 features used to describe users of the service in their tweets. We included within our contextual features, anything that told us about the user that was freely provided. This ranged from the sensed (for example their location details) to the surveyed (their Twitter biography), all accounting for the context of how that user presents themselves to others. We took 37 features made available in the tweets (such as the source and which software client sent the tweet) or otherwise computable from the features available. These features constituted the contextual attributes of a user, visible to others, making them factors which could impact friend recommendation. Where we knew the feature would be unique (such as the screen name or real name of a person) we computed features that would make these fields comparable (detailed in the list below). In addition to these 37 features we had 24 features to characterise how many times the user tweets in each hour of the day. For the pur-

[3]http://twitter.com

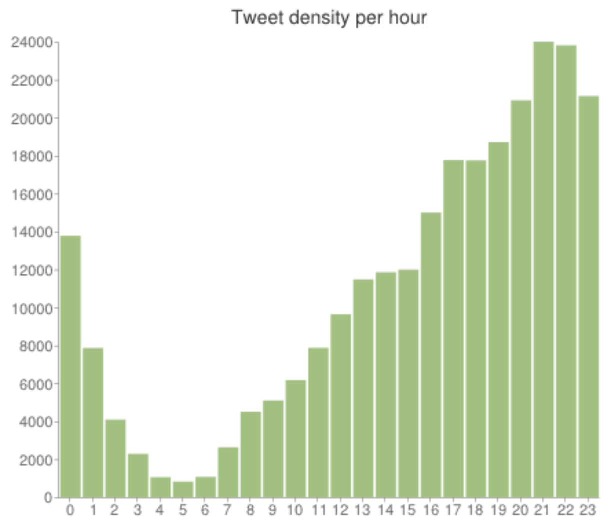

Figure 1: Tweet density over time, from public Dublin-based Twitter users over time

poses of using machine learning we categorised each of the text features with a number, Table 1 details the number of categories generated for each text feature. This preprocessing gave us a list of 7,390 users as described by the context they present to the world, that they tweet only at certain times, or are popular or unpopular (based on follower count or similar metrics). We computed the following categories to represent attributes that could offer other insight into following habits, detailed in the list below.

Capital letters in screen name Number of capital letters in user's screen nickname.

Capital letters in name Number of capital letters in user's actual name.

Description length Number of characters in the user's biographical description.

Name length Number of characters in the user's name

Screen name length Number of characters in the user's screen name

Screen name is real name Is the user's screen name equivalent to their real name?

For the purposes of our experiment we were interested in who each user in the collection followed, and what contextual data might have influenced that decision. We gathered each person in the collection's complete friends list. This allowed us to highlight which people in the collection followed each other. We were then able to generate for each person, a list of every other user in the collection as described by their contextual features, annotated with whether or not that person follows them. This pre-processing left us with the data formatted for the tests we wished to perform. In our analysis we first looked at the importance of each feature as a means of discriminating within the set for each user. F-score is a simple technique which measures the discrimination of two sets of real numbers, as described in [27]. The larger the F-score is, the more likely this feature is to

Table 1: Text features and the number of categories for each

Feature	Number of Categories
Geotype	1
language	11
location	2552
place full name	38
place id	38
place name	35
place type	3
place URL	35
prof back colour	1089
prof sidebar colour	1116
prof sidebar fill colour	1180
prof text colour	1021
source	101
timezone	75

be more discriminative. It is important to note that if one user exclusively follows people with low tweet counts and another exclusively follows people with high tweet counts then both will have high F-scores, as "number of tweets" is a very discriminative feature for both. We calculated the importance of each feature for every user, then averaged them over all users. This will form an integral part of the feature selection we perform later. For each person within the set we computed their individual F-scores based on who they followed.

Having examined F-scores, we then proceeded to perform feature selection for a group of 530 users from the collection. We did this in order to see what influenced whether one person followed another, in order to potentially offer better contextual recommendation. We used this data to build an SVM for each person, using libSVM ([7][4]). Training used the entire list of users with the 61 features and whether or not the user in question follows them. We categorised all of the text-based features into numerical format in order to be compatible with the SVM training. We used the feature selection tool provided with libSVM[5] to rank the important features in the dataset. Afterwards we ran feature selection on each user, and then we took the minimum number of features necessary to accurately produce the same results in order to arrive at our final analysis.

4. RESULTS

Examining the F-scores of each contextual feature we found, as detailed in Table 2, that each of the most important features has a high standard deviation, indicating importance of features is very personal to each user. We see that on average, follower count is clearly the most discriminating feature. Examining the figures for the tope five features, followers, listed, friends, favourites and statuses (numbered 1 to 5 respectively in 2), as a boxplot makes this even more clear; there is a contextual long tail of users who place varying degrees of importance on the features outside the 25%-75% quartile box.

Having looked at the most discriminating features available, we then trained 530 individual SVMs, one for each user.

[4]http://www.csie.ntu.edu.tw/ cjlin/libsvm/

[5]http://www.csie.ntu.edu.tw/ cjlin/libsvmtools/

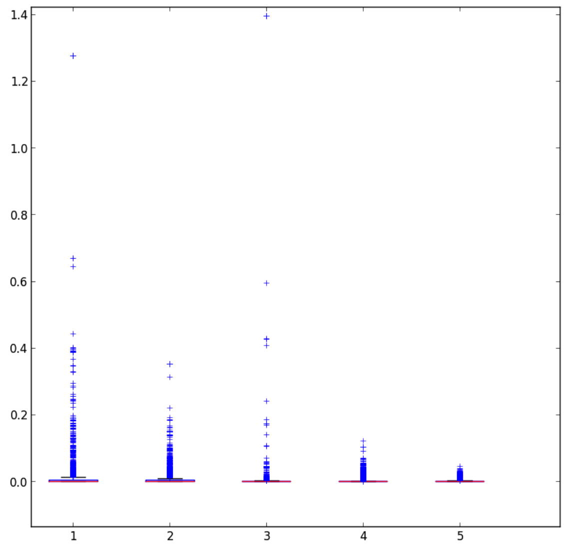

Figure 2: The difference in effect of contexts for individuals is clear.

Table 2: The top average important features in deciding whether a user follows another

Feature	Strength	Std Dev
Follower count	0.01147	0.0689
Listed count	0.00673	0.0236
Friends count	0.00260	0.0402
Favourites count	0.00243	0.0077
Statuses count	0.00147	0.0037
Posts during 16:00	0.00093	0.0070
Posts during 19:00	0.00068	0.0048
Posts during 17:00	0.00067	0.0039
Posts during 21:00	0.00065	0.0038
Posts during 20:00	0.00065	0.0036
Posts during 22:00	0.00065	0.0035

These SVMs were trained on the prepared list of each users' contextual representation, annotated by whether or not the SVM-focused user follows them. In all but three cases, users' following habits were indicated by only three features. The three special cases include one user who required 13 features and two that maintain their highest accuracy with six features. This shows that anTable 3 shows an aggregated count of features as they appear across each user's feature selection set. This corresponds to how the user evaluates who they follow, meaning that (all other things being equal) a highly discriminating feature helps easily decide whether or not to follow. Follower count and Listed count, both highly discriminating features overall, top the list, but other features that may not be as obvious, such as whether the profile background of a user is tiled, play a part in defining how users choose who they follow.

5. CONCLUSIONS

5.1 Discussion

Table 3: The most selected features by SVMs trained on individual users

Feature	Users
Follower count	185
Listed count	174
Profile background is tiled	90
Description length	81
Statuses count	64
Screen name is real name	63
Geotype	58
Favourites count	53
Name length	51
Profile text colour	49
Friends count	47
Location	39
Capital letters in screenname	33
Capital letters in name	32
Profile sidebar border colour	30
Posts during 12:00	29
Source	29
Profile background colour	27
Posts during 14:00	23
Place name	22
Posts during 7:00	21

On the social web people use context in somewhat unique ways. While the posting habits of a potential friend suggestion matters to some people it, as we have shown, does not matter to others. We have shown here that there are groups of users who use entirely different sets of contexts. Of the 530 users we examined no feature was globally selected as a positive predictor, in fact the best predictor was only commonly selected for 35% of users, showing that no single contextual factor can be chosen to optimally improve recommendation. There is, on an individual level, a high degree of variation when it comes to how important users find this information. In effect, people have different priorities around contexts, and we have outlined a way to detect and highlight them on an individual level. There is a long tail of users who place greater or lesser importance on every contextual information.

Depending on the user, we can recommend the *context set* they should use, in order to improve contextual recommendation in the case of recommendations of whom a user should follow in Twitter. This could conceivably lead to modelling users based on what criteria they use to evaluate the world, a "context-profile" that could accompany people in the cloud to be used by any service that recommends using context. We have highlighted in Twitter that follower-count is a decisive metric for user interest in following. However it is only seen as important to 185 of the 530 people who were analysed, indicating that it would not improve recommendation for the majority of users. If there had been some solid consensus on which features to use this would be a valid method of using context to choose the context to use when recommending. If further investigation found this to be a wholly positive correlation (i.e. people always valued more followers) this could speak to the suitability of collaborative filtering for Twitter user recommendation, as sparse ratings would actually be indicative of trend toward a less suitable recommendation.

Furthermore, it is interesting to note that while no contextual feature provides good coverage of the 530 users (i.e. no one feature could be used to predict accurately), sets of contexts do reoccur, opening up the possibility of using a recommender system to class users based on their behaviour and recommend a set of contexts that will most likely improve their recommendations.

This work shows that context has a variable effect on the decision making of users of the social web, and any recommender system that hopes to incorporate the wealth of data available from social networks should be aware of this detectable difference.

5.2 Future Work

This research shows that people on the social web have differing priorities in terms of context, and we can detect them. This detection could be used to create an ontological model of the importance users place on different sensed contexts, effectively allowing their attitudes toward contexts choose which factors influence their recommendations. This raises further questions beyond the scope of this paper, which we hope to study in our ongoing research. Firstly we wish to quantitatively analyse the impact on recommendation accuracy the methods of pre-filtering/post-filtering and user modelling only important context rather than all available contexts, and possible context selection strategies that arise from this.

Memory-based recommenders could make use of these ontological models in grouping. Rather than grouping users because they both like similar things in similar situations they could be grouped based on similar items, rated in different situations, because they both have the same ontological interests in context. This is one area we wish to explore in more detail. Further this sort of examination can evaluate whether support for contextual features, which can be costly in either sensor deployment or user effort, are actually making a difference, and optimise sensor sets for a range of users on specific recommendation tasks.

Acknowledgments

The authors acknowledge support of Science Foundation Ireland through grant number 07/CE/I1147.

6. REFERENCES

[1] G. Adomavicius and A. Tuzhilin. Context-aware recommender systems. In F. Ricci, L. Rokach, B. Shapira, and P. B. Kantor, editors, *Recommender Systems Handbook*, pages 217–253. Springer US, 2011.

[2] S. S. Anand, P. Kearney, and M. Shapcott. Generating semantically enriched user profiles for web personalization. *ACM Trans. Internet Technol.*, 7(4), Oct. 2007.

[3] L. Baltrunas, B. Ludwig, and F. Ricci. Context relevance assessment for recommender systems. In *Proceedings of the 16th international conference on Intelligent user interfaces*, IUI '11, pages 287–290, New York, NY, USA, 2011. ACM.

[4] J. S. Breese, D. Heckerman, and C. M. Kadie. Empirical analysis of predictive algorithms for collaborative filtering. In G. F. Cooper and S. Moral, editors, *UAI*, pages 43–52. Morgan Kaufmann, 1998.

[5] M. Brunato and R. Battiti. Pilgrim: a location broker and mobility-aware recommendation system. In *in Proceedings of IEEE PerCom2003*, page submitted. IEEE Computer Society, 2002.

[6] M. Brunato and R. Battiti. PILGRIM: A location broker and mobility-aware recommendation system. In *Pervasive Computing and Communications, 2003. (PerCom 2003). Proceedings of the First IEEE International Conference on*, pages 265–272, March 2003.

[7] C.-C. Chang and C.-J. Lin. Libsvm: A library for support vector machines. *ACM Trans. Intell. Syst. Technol.*, 2(3):27:1–27:27, May 2011.

[8] L. S. Connaway, T. J. Dickey, and M. L. Radford. âĂİ If it is too inconvenient I'm not going after it:âĂİ Convenience as a critical factor in information-seeking behaviors. *Library & Information Science Research*, 33(3):179–190, July 2011.

[9] J. Derrida. *Of grammatology / by Jacques Derrida ; translated by Gayatri Chakravorty Spivak.* Johns Hopkins University Press, Baltimore, 1st American ed. edition, 1976.

[10] A. K. Dey. Understanding and using context. *Personal and Ubiquitous Computing*, pages 4–7, 2001. http://www.odysci.com/article/1010112991760792.

[11] R. Dhar, S. Nowlis, and S. Sherman. Trying hard or hardly trying: An analysis of context effects in choice. *Journal of Consumer Psychology*, 9(4):189–200, 2000.

[12] C. Goodwin and A. Duranti. *Rethinking context : language as an interactive phenomenon / edited by Alessandro Duranti and Charles Goodwin.* Cambridge University Press, Cambridge [England] ; New York :, 1992.

[13] Google, Ipsos, and OTX. The Smart Shopper, Sep 2010.

[14] P. Ingwersen. Cognitive perspectives of information retrieval interaction: elements of a cognitive ir theory. *Journal of documentation*, 52(1):3–50, 1996.

[15] P. Ingwersen and K. Järvelin. *The Turn: Integration of Information Seeking and Retrieval in Context (The Information Retrieval Series).* Springer, 2005.

[16] Y. Koren. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, KDD '08, pages 426–434, New York, NY, USA, 2008. ACM.

[17] O. Madani and D. DeCoste. Contextual recommender problems. In *Proceedings of the 1st international workshop on Utility-based data mining*, UBDM '05, pages 86–89, New York, NY, USA, 2005. ACM.

[18] W. Newman and W. M. Newman. Better or just different? on the benefits of designing interactive systems in terms of critical parameters. In *In Designing Interactive Systems (DIS97*, pages 239–245. ACM Press, 1997.

[19] A. Noulas, S. Scellato, N. Lathia, and C. Mascolo. A random walk around the city: New venue recommendation in location-based social networks. In *IEEE International Conference on Social Computing*, 2012.

[20] K. Oku, S. Nakajima, J. Miyazaki, and S. Uemura. Context-aware svm for context-dependent information

recommendation. In *Mobile Data Management, 2006. MDM 2006. 7th International Conference on*, page 109, may 2006.

[21] U. Panniello, A. Tuzhilin, M. Gorgoglione, C. Palmisano, and A. Pedone. Experimental comparison of pre- vs. post-filtering approaches in context-aware recommender systems. In *Proceedings of the third ACM conference on Recommender systems*, RecSys '09, pages 265–268, New York, NY, USA, 2009. ACM.

[22] M. Park, J. Hong, and S. Cho. Location-based recommendation system using bayesian userâĂŹs preference model in mobile devices. *Ubiquitous Intelligence and Computing*, pages 1130–1139, 2007.

[23] R. Savolainen. Everyday life information seeking. *Encyclopedia of Library and Information Science*, pages 1780–1789, 2010.

[24] A. Schmidt, M. Beigl, and H. w. Gellersen. There is more to context than location. *Computers and Graphics*, 23:893–901, 1998.

[25] E. Shen, H. Lieberman, and F. Lam. What am i gonna wear?: scenario-oriented recommendation. In *Proceedings of the 12th international conference on Intelligent user interfaces*, IUI '07, pages 365–368, New York, NY, USA, 2007. ACM.

[26] M. Warner. Publics and counterpublics. *Chicago: Duke University Press*, 2005.

[27] Y. wei Chen and C. jen Lin. Combining svms with various feature selection strategies. In *Taiwan University*. Springer-Verlag, 2005.

[28] T. Wilson. Models in information behaviour research. *Journal of Documentation*, 55:249 – 270, 1999.

[29] W.-S. Yang, H.-C. Cheng, and J.-B. Dia. A location-aware recommender system for mobile shopping environments. *Expert Systems with Applications*, 34(1):437 – 445, 2008.

[30] S. O. Yoon and I. Simonson. Choice set configuration as a determinant of preference attribution and strength. *Journal of Consumer Research*, 35(2):324–336, 2008.