Co-Following on Twitter

Venkata Rama Kiran Garimella*
Aalto University
Espoo, Finland
kiran.garimella@aalto.fi

Ingmar Weber
Qatar Computing Research Institute
Doha, Qatar
iweber@qf.org.qa

ABSTRACT

We present an in-depth study of co-following on Twitter based on the observation that two Twitter users whose followers have similar friends are also similar, even though they might not share any direct links or a single mutual follower. We show how this observation contributes to (i) a better understanding of language-agnostic user classification on Twitter, (ii) eliciting opportunities for Computational Social Science, and (iii) improving online marketing by identifying cross-selling opportunities.

We start with a machine learning problem of predicting a user's preference among two alternative choices of Twitter friends. We show that co-following information provides strong signals for diverse classification tasks and that these signals persist even when the most discriminative features are removed.

Going beyond mere classification performance optimization, we present applications of our methodology to Computational Social Science. Here we confirm stereotypes such as that the country singer Kenny Chesney (@kennychesney) is more popular among @GOP followers, whereas Lady Gaga (@ladygaga) enjoys more support from @TheDemocrats followers.

In the domain of marketing we give evidence that celebrity endorsement is reflected in co-following and we demonstrate how our methodology can be used to reveal the audience similarities between not so obvious entites such as Apple and Puma.

1. INTRODUCTION

How much does following a particular set of people reveal about your interests? Does the fact that you follow @Starbucks make it more likely that you follow @TheDemocrats as well? And can Twitter users be grouped in a meaningful way by looking at whether their followers have similar friends[1]?

Such questions are relevant to at least three lines of research. First, there is lots of work on user classification on Twitter, e.g., [7, 8, 3]. Such classifiers often rely on language-specific tools such as stemming or dictionaries with special terms. Our work shows that such information might not be required and a language-agnostic method using a user's friends as features achieves ROC-AUC of .80-.85 for a wide range of binary classification tasks. Second, online social networks are becoming a more and more important data source for Computational Social Science [16, 6]. We contribute to this area by showing how things such as "lifestyle politics" can be studied by using co-following information. Lastly, Twitter with its user base of several hundreds of millions is an important advertising and marketing platform. We show how followership-based similarity methods can be used to identify accounts with a similar audience in terms of interests which could create cross-selling opportunities.

Most of our analysis is centered around 18 rivalries such as @GOP vs. @TheDemocrats or @McDonalds vs. @BurgerKing. In many cases the two alternatives are arguably interchangeable and one might not expect a big difference between the interests of the followers of, say, @Hertz and @Avis. For each of these seed pairs we obtained up to 2,000 random followers. Their friends are used to construct feature vectors and we perform an in-depth analysis of the co-following behavior. We also construct the same kind of vectors for a set of popular musicians, and in all cases the basic hypothesis is that users following similar users are similar and that this propagates to their friends. Our findings include the following.

1. Using *solely* language-agnostic co-following information provides strong signals concerning a user's preference even among arguably interchangeable choices such as @Hertz or @Avis.

2. Such classification is robust with respect to the removal of the most strongly and often obviously related co-following features.

3. Aggregated signals from the general crowd work better for distinguishing binary preferences than relying on the most similar users in a k-NN fashion.

4. A feature analysis confirms stereotypes such that @ladygaga is more popular among @TheDemocrats followers, but also reveals less expected patterns such as that @SnoopDogg followers tend to prefer @Pepsi over @CocaCola.

5. There is evidence that celebrity endorsement works as following a celebrity increases the probability of following the related product.

6. Groups of related Twitter accounts, such as musicians, can be mapped in a simple manner by looking at their followers' friends.

7. Such a mapping reveals Republicans' preference for Miller's beers, Democrats' preference for Budweiser, and the fact that Apple and Puma target a similar, metropolitan audience.

To the best of our knowledge, this is the first such study of co-following on Twitter. We hope that both our analysis and our tools will be of interest to researchers working on user classification, Computational Social Science, or on social media marketing.

*Most of this work was done while the author was at QCRI.

[1]We use the term "friend" as Twitter terminology referring to another Twitter user that a user follows.

HT '14, September 01 - 04 2014, Santiago, Chile.

http://dx.doi.org/10.1145/2631775.2631820.

2. RELATED WORK

A key assumption to group users based on the similarity of their followers' friends is that following is an expression of topical interests or demographic similarity, rather than personal contacts. Kwak et al. [15] observe that Twitter is more of a news source than a social network, which is *good* for our applications as this indicates that follower/friend links are more related to user interests than social connections.

Related to our approach of using co-following to measure user similarity is the work in [2] where An et al. use the common audience of two Twitter accounts as a measure of closeness. This differs fundamentally from our methodology as we use *second order* co-following. Concretely, two accounts that do not share a single follower are considered similar by us if their followers share many friends. We believe that such an approach is preferable to break out of the "homophily ghetto" where users follow what their friends follow. It also allows for a more far-reaching notion of similarity that could be used to, say, align political parties in different countries on a common spectrum, even if no user follows parties in two distinct countries [24].

Several papers have looked at various Twitter user classification tasks, typically for (i) political orientation in the US, (ii) gender, and (iii) age [19, 21, 26, 7, 8, 3]. This line of work usually involves a broad set of features, including textual content, network and activity based features, as well as a variety of classification approaches that make use of label-propagation across social links. Our approach differs from these in a number of ways. First, the binary classification task of "does the user follow A or B" is different. Second, we do not use any content-based features. Third, we do not use any retweet signals as we are not interested in the sparse network of interests that users strongly engage with, but rather the larger network of "weak interests". Fourth, we do not make use of label-propagation across social links as we are not interested in methods that work for (and re-inforce) "information bubbles" but we are looking for approaches that can be transferred to completely new domains where users do not yet have any direct social ties. Finally, the actual classification performance is of less interest to us than an understanding of how much information is contained in co-following and how this could be used for different applications.

Our mapping and visualization of similar accounts is conceptually similar to community detection/clustering which also identifies groups of related accounts [17, 18, 11]. Our approach, presented in Section 5 is different from these, as we do not require a global view of the entire network as we are only interested in understing the relative positions of the main users. Also, we do not want to find communities induced by friends-of-friends type links. Rather we strive for a similarity-only based approach that can easily be transferred to domains without any friends-of-friend links.

Determining which of two alternatives a Twitter user is more likely to follow is related to friend recommendation or link prediction as, in a sense, we are suggesting which of the two links should be formed. Intuitively, transitivity and mutuality of links are important signals for link prediction [12] but, as discussed previously, we do not want to use such "three people you follow also follow X" information as it leads to a different type of application, closer to community detection. User similarity based on user attributes has also been used as a feature for link prediction [25, 13]. But this work still partly relies on mutuality and transitivity, which is equivalent to re-enforcing partisan camps without noting any existing similarities in terms of shared interests. For example, such approaches would most likely fail to pick up the similarity between @Puma and @TheAppleInc that we observe. Weber et al. [24] present applications of the idea of second order co-following for making out-of-context recommendations of musicians and politicians.

3. DATA

Our data set is constructed around a set of *Twitter seed accounts*. These accounts correspond to "rivalries" between two entities such as @CocaCola vs. @Pepsi or @Samsung vs. @TheAppleInc. The full list of 18 account pairs can be found in Table 1. The list of these rivalries was obtained from Fortune.com.[2] Later, we also look at *groups* of seed accounts, namely, Twitter accounts for (i) popular musicians, and (ii) all the 18 rivalries combined. In all cases, we first obtained a list of all the accounts' followers. From this list we then sampled uniformly at random a set of 2,000 followers. For each of the sampled followers we obtained the full list of their Twitter friends, i.e., users that they follow. The sampling of 2,000 followers was done in order to make this step of acquiring the friends feasible (due to the strict rate limits of the Twitter API). The seed accounts pertaining to the corresponding rivalry/group were removed from these friends lists and the remaining ones were treated as a feature vector with each dimension corresponding to a Twitter account being followed. Users who followed *only* seed accounts were dropped. For the cases of rivalries, we also imposed the constraints that the followers were located in the United States. This was done to avoid picking up differences in international market penetrations, rather than within-US cultural differences.

We used this data to construct a binary classifier for which we created train and test splits, each consisting of $\sim$1,000 users. Note that even though 2,000 users were sampled, due to limitations in the Twitter API, the actual number of users for which we could get the friends varies between 1,800-2,000. This could be due to changes in users' privacy settings, accounts getting blocked and so on. For constructing the training vectors, we only considered users who were followed by at least two users in our training set. That is, if only one of the thousands of users in the training set followed @phdcomics, then following @phdcomics would not be used as a feature. This serves as a simple method for reducing the dimensionality as well as removing unimportant dimensions. This is analogous to the text mining scenario of removing rare tokens with a frequency 1.

4. CO-FOLLOWING AND BINARY PREFERENCES

In this section we look at how much a user's choice of Twitter friends reveals about their preference among two alternatives such as @CocaCola vs. @Pepsi. We do this with different research questions in mind. First, we approach things from a machine learning perspective with an evaluation of the corresponding binary classification task. Next, we do a feature analysis to see which arguably irrelevant features, such as musicians followed, provide information about a user's soft drink or political preferences.

Feature Vectors Using IDF. As a preprocessing step, we transformed our binary X-follows-Y vectors to an IDF-weighted alternative. To illustrate why, imagine that almost everybody follows @SuperCelebrity. Then following @SuperCelebrity is not very informative or discriminative and is given a very low IDF weight. For each of the 18 rivalries, we compute the IDF scores of the friends of the followers of the seed 36 seed rivals. In total, 63,853 (N) followers of the seed accounts were used (= 36 x 2,000, minus cases with fewer than 2,000 followers and blocked/deleted/private accounts).

[2] `http://bit.ly/1o6WMqf`

We computed the IDF of each of their friends, obtaining one IDF-weighted vector for each of the 63,853 followers.

$$IDF(user_i) = log\left(\frac{N}{|followers(user_i)|}\right), \qquad (1)$$

where $followers(user_i)$ indicates the followers of a particular user, from the set of followers sampled for the seed rival accounts. Each of these IDF-weighted vectors was then normalized in 2-norm and, for a given seed account, all of its followers normalized vectors are summed up. This summed vector is then re-normalized in 2-norm to give the final "global" summary vector for the seed account.

4.1 Machine Learning Performance

In this section we evaluate how much information co-following provides for the task of classifying users according to their binary preference. The ground truth is the single account that the test user actually followed. The feature dimensions corresponding to the following of the seed accounts are always removed and, later, we also remove strongly correlated features such as following @BarackObama for the @GOP vs. @TheDemocrats task. Empty vectors, after removing the seed accounts (for users following only the seed accounts) are ignored. Our main performance measure is the area-under-curve (AUC) for the Response Operating Characteristic Curve (ROC) as computed by using a toolkit provided by Goadrich, et al. [9]. We also report AUC for the Precision-Recall Curve (PR) though AUC-ROC will be the default. A value of 0.5 indicates a random, unskilled prediction model.

Note that we are more interested in understanding the relative performance when, say, the most discriminative features are removed than we are in achieving the highest possible classification accuracy. The accuracy could always be improved further by using other algorithms (SVM, Maximum Entropy, etc.), other feature sets (textual data, interaction features, network features, etc.), or incorporating other techniques (label propagation, community detection, etc.). Our focus is more on understanding issues such as robustness under feature removal, relative performance on sparse test vectors or opportunities for Computational Social Science arising from feature analysis.

Global vs. Local Approach. To determine whether you fall into group A or B, is it more useful to know (i) what the general, average members of A and B are like, or (ii) which if any of the two contains a small number of members just like you? The answer to this question has applications both for the design of classification algorithms and for understanding the structure of groups of followers. We try to answer this question by comparing two different classification strategies. First, a "global" method using the single IDF-weighted summary vector described above. This method also includes information about fairly rare friends as it aggregates information from about 2,000 followers. Second, a "local" approach, that uses a k-nearest neighbor classifier. It then assigns each test vector to the class with the largest number of close neighbors among the top k. For k-NN we experimented with a range of values for k from 1 to 9 in increments of 2. There was a clear tendency for higher values of k to perform better and so we stuck to a choice of k=9. We did not experiment with larger values as the general trend of a more and more global approach performing better was our main objective, rather than identifying an optimal value of, say, k=135.

The performance of the binary classification is shown in Table 1. The global approach always performs better than the local approach, showing that the it is worth aggregating the long tail of rare co-follower relations. The AUC-ROC averaged across the 18 tasks is 0.81.

Rivalry	Global	Local
@Budweiser vs. @MillerCoors	0.86 (0.91)	0.80 (0.85)
@FedEx vs. @UPS	0.73 (0.73)	0.69 (0.72)
@GM vs. @Ford	0.75 (0.86)	0.69 (0.76)
@GOP vs. @TheDemocrats	0.91 (0.95)	0.86 (0.93)
@Hertz vs. @Avis	0.92 (0.93)	0.91 (0.92)
@InsideFerrari vs. @lamborghini	0.92 (0.95)	0.87 (0.93)
@jcpenney vs. @Sears	0.75 (0.82)	0.67 (0.72)
@McDonalds vs. @BurgerKing	0.78 (0.79)	0.68 (0.70)
@MercedesBenz vs. @bmw	0.89 (0.93)	0.86 (0.91)
@Nike vs. @Reebok	0.78 (0.74)	0.73 (0.68)
@NikonUSA vs. @CanonUSAimaging	0.83 (0.85)	0.78 (0.83)
@pepsi vs. @CocaCola	0.69 (0.76)	0.65 (0.73)
@PUMA vs. @adidas	0.77 (0.84)	0.69 (0.73)
@SamsungMobile vs. @TheAppleInc	0.95 (0.96)	0.92 (0.94)
@Starbucks vs. @DunkinDonuts	0.80 (0.87)	0.72 (0.82)
@Target vs. @Walmart	0.78 (0.86)	0.69 (0.79)
@thewanted vs. @onedirection	0.79 (0.88)	0.76 (0.84)
@Visa vs. @MasterCard	0.71 (0.72)	0.62 (0.59)

Table 1: Performance comparison for the 18 binary classification tasks (detecting preference among rivaling alternatives) in terms of AUC-ROC (AUC-PR) for both the global and local similarity-based approaches.

Removing Obvious Co-following Signals. Discovering that following @BarackObama on Twitter is an indication for following @TheDemocrats rather than @GOP is obvious. Similarly, following @CokeZero correlates positively with following @CocaCola. As we were more interested in studying the *non-obvious* dependencies we investigated the classification performance when the most predictive features are removed. Note that this is the *opposite* of what normal feature selection does.

Concretely, for each binary setting we rank features as follows. For each rivalry pair A, B, we look at the absolute differences in the feature values listed in A's and B's summary vectors. These absolute differences are then sorted in descending order. Features with a large difference correspond to accounts that are typically more followed by the followers of one seed account, but not the other one.

In order to check the influence of the top features, we removed the top 10, 20, 50, 100 and 200 most obvious features and compared the AUC in each case. Note that in this setting, since we remove the most influential features, the size of the test set might change (because some users might only follow these influential users). In order to compensate for this, we tried two variants, one considering only users who have more than 201 followers, so that the size of the test set is fixed and the other with a varying test set size. The results of the former case are presented in Figure 1, though results in both cases are comparable. The y-axis indicates AUC averaged across all the rival groups. We note a gradual decrease in the mean AUC as we remove more features, which is in line with what is expected.

For our later "mapping" analysis (Section 5), we take a similar approach to remove the top 20 features for each of the A vs. non-A classification problems, where A iterates over all the seed accounts and the non-A group pools all non-A seeds.

4.2 Feature Analysis

Correlation with Lifestyle. Apart from using it for user classification and targeted advertising, co-following patterns are also of interest in their own right and can serve to answer questions in Computational Social Science. For instance, there is academic work that

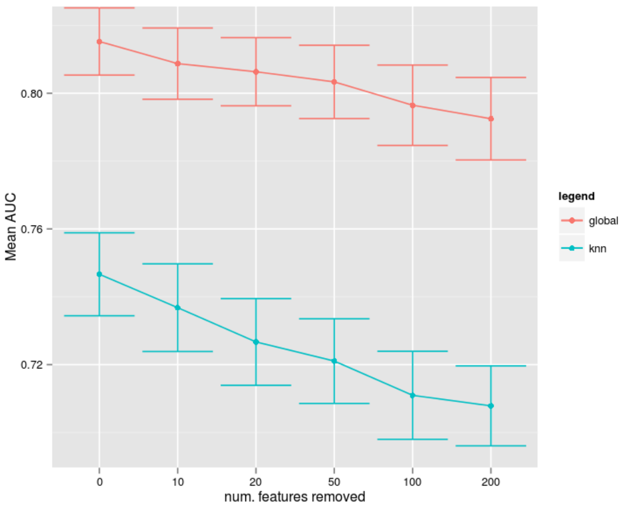

Figure 1: Average AUC across the 18 binary classification tasks (detecting preference among rivaling alternatives) as more and more features are removed, but the test set size is fixed. Only users with more than 201 followers are considered. Error bars indicate the standard error across the tasks.

looks at "lifestyle politics" such as the correlation between political leaning and television preferences or the stereotype that liberals like lattes [5, 4, 20]. Our approach contributes to this by offering a language-agnostic tool to use online data to quantify such effects at scale. In this section, we use the feature ranking described previously to generate the top, discriminative features (such as following @BarackObama to predict a @GOP or @TheDemocrats preference).

As an example, we look at the rivalries between @TheDemocrats vs. @GOP, and @Pepsi vs. @Cocacola. For both cases we look at the top discriminative co-following features from the `http://WeFollow.com` categories Music, Sports and News.[3] Table 2 shows some examples of the insights we can get from co-following patterns. The lifestyle correlations for the political rivalry @GOP vs. @TheDemocrats can be inspected to make intuitive sense with, e.g., @nytimes being more popular among @TheDemocrats followers.[4] For @Pepsi vs. @CocaCola many observations can be explained by the fact that @Pepsi targets the younger "New Generation".

Celebrity Endorsements. An interesting side note of examining the top features is the detection of celebrity endorsements. By observing these features, we found out that celebrity endorsements go hand in hand with who people follow. Some examples include Derrick Rose (@drose) for Adidas; Kevin Durant (@KDTrey5) for Nike; SnoopDogg (@SnoopDogg), Nicki Minaj (@NickiMinaj) and Drake (@drake) for Pepsi, and Maroon 5 (@maroon5) and David Guetta (@davidguetta) for Cocacola, etc. This observation has interesting applications in marketing campaigns and recommendation systems and deserves more analysis in the future.

5. MAPPING THE TWITTERSPHERE VIA CO-FOLLOWING

[3]A small number of mislabeled entries were removed. For example, @AgainAmerica was incorrectly listed in the Music category.

[4]The New York Times is generally perceived to have a liberal bias, see `http://en.wikipedia.org/wiki/The_New_York_Times#Political_persuasion_overall`.

	@GOP	@TheDemocrats
Music 1	@kennychesney (64)	@ladygaga (24)
Music 2	@jakeowen (122)	@aliciakeys (57)
Music 3	@taylorswift13 (139)	@SnoopDogg (62)
Sports 1	@espn (125)	@rolandsmartin (66)
Sports 2	@runnersworld (143)	@bubbawatson (79)
Sports 3	@AdamSchefter (178)	@NBA (188)
News 1	@WSJ (12)	@nytimes (11)
News 2	@HumanEvents (77)	@cnnbrk (21)
News 3	@toddstarnes (107)	@NYTimeskrugman (23)
	@Pepsi	@Cocacola
Music 1	@SnoopDogg (1)	@maroon5 (8)
Music 2	@Nickiminaj (2)	@davidguetta (19)
Music 3	@Drake (5)	@Pitbull (28)
Sports 1	@shaq (3)	@SInow (99)
Sports 2	@ochocinco (20)	@kaka (103)
Sports 3	@DwightHoward (42)	@chicagobulls (206)
News 1	@Rapup (120)	@cnnbrk (55)
News 2	@Life (133)	@WSJ (81)
News 3	@MTVnews (339)	@TheOnion (183)

Table 2: List of differentiating co-following features from different WeFollow classes, for @GOP vs. @TheDemocrats, and @Pepsi vs. @Cocacola. The numbers in parentheses indicate the absolute position of this feature in our ranking irrespective of the topic (Music, Sports or News).

In this section we look at whether a co-following based similarity can be used to map the relative positions of players from domains such as music. Though our maps can be seen as "community detection", the approach and interpretation is very different. Whereas traditional community detection algorithms use direct social links and, e.g., would try to find clusters with unusually high triadic closure [17], our approach relies on more indirect and high-order links. As an example, imagine two football clubs that are fierce rivals and who would definitely not follow each other. Fans and Twitter followers of either club might also not follow the other one. However, their fans might jointly follow many other accounts related to sports news. Due to this co-following of the clubs' followers we would consider the clubs as similar in terms of their audiences' interests. This approach also opens up opportunities for cross-marketing and cross-selling: if two Twitter accounts from different domains share a similar followership then they might consider cross-posting or otherwise combining their forces. Note again that they do not have to share even a single follower to be considered similar as we look at second-order following relations, namely, the friends of their followers.

Technically, we did the following: For each of the ~2,000 followers of an account, we constructed the IDF vector for the users they follow.[5] We then computed the pair-wise cosine similarity between these feature vectors. Since we need distances, we used (1 - cosine similarity) as the measure of distance. We then used the classical, Metric Multi-Dimensional Scaling [23] (MDS) on this data with the $cmdscale()$ function in $Matlab$. Note that MDS is a *lossy* embedding and that even though two points appear close in the 2-dimensional plane, they might be far apart in the original high dimensional space. Therefore, all conclusions and observations we derived from such mappings in the following have also been validated using the high dimensional similarity information.

[5]The total number of followers (N) for Musicians was 38,358.

5.1 Popular Musicians

To see if our approach generalizes to diverse domains such as music, we decided to map popular musicians on Twitter. To this end we obtained a list of the top 22 musicians from `http://wefollow.com/interest/music`.[6]

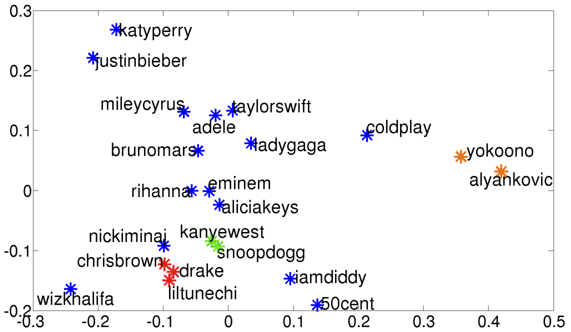

Figure 2: A 2D MDS similarity map of popular musicians. Similarity measures are derived from their followers' aggregated friends.

Figure 2 shows the map that was obtained using MDS on the musicians data. Most of the observed structure corresponds to musical genres. For example, Lil Wayne (@liltunechi), Chris Brown (@chrisbrown) and Drake (@drake) are rappers and are co-mapped together in the map, marked in red. Similar is the case of Snoop Dogg (@snoopdogg) and Kanye West (@kanyewest), marked in green, both of which are hip hop artists. However, there are also surprising things that emerge such as the relative closeness of "Weird Al" Yankovic (@alyankovic), famous for musical parody, and Yoko Ono (@yokoono), both marked in orange. Though very different musical genres, both arguably appeal to an older, more educated audience. This already hints at applications of such analysis for the identification of cross-selling opportunities.

5.2 Combination of All Rivals

To show the full generalizability of this mapping approach also *across* domains we combined all the 36 Twitter accounts from the 18 rivalries and mapped them in a common space in Figure 3. As one would expect, many rivals such as @Target vs. @jcpenney and @thewanted vs. @onedirection are comparatively close as their followers share similar interests. However, the relative distances across rivalries also makes sense. For example, the beer brand @MillerCoors is closer to @GOP than to @TheDemocrats and the opposite holds for @Budweiser. This makes sense as it has been observed before that "Republicans are also big fans of Miller Lite and Coors Light, but Democrats drink more Budweiser" [10, 22], though, some studies show that the opposite is true [1]. Sometimes, studies such as this one are inconclusive and show conflicting results based on the demographics studied (such as voters vs. just politically leaning but not necessarily voting), sampling methods used, etc. Similarly, @GM and @Avis are very close in the low-dimensional embedding. Again, this makes sense as "[s]ince the late 1970s, Avis has featured mainly General Motors (GM) vehicles"[7]. Also noteworthy is the closeness between @PUMA and @TheAppleInc. Though there does not appear to be any formal alliance, both brands try to create a similar image of themselves. Puma targets the "sports lifestyle" trend with persona attributes such as metropolitan and international [14] which, arguably also applies to Apple. We believe that such a mapping is useful to quickly generate hypotheses for lifestyle politics and similar research areas that can then be investigated in depth. It is important to note that even if some of these findings were not to hold "offline" in all cases, these Twitter-only findings are still useful for online advertising as they definitely provide a signal.

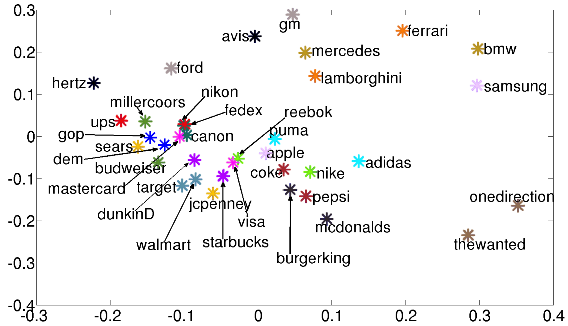

Figure 3: A 2D MDS similarity map of all the 18x2 rivals. Similarity measures are derived from their followers' aggregated friends. Rival pairs are represented by stars of the same color. Some labels have been shortened due to space constraints. dunkinD represents DunkinDonuts and dem., Democrats.

6. CONCLUSIONS

We presented an in-depth study of co-following behavior on Twitter which contributes to (i) a better understanding of language-agnostic user classification on Twitter, (ii) eliciting opportunities for Computational Social Science, and (iii) improving online marketing by identifying cross-selling opportunities. Concretely, we used the similarity of followers' friends to predict a users' preferences and to group main Twitter users according to their audiences' similarities. We showed that such language-agnostic co-following information provides strong signals for diverse classification tasks and that these signals persist even when the most discriminative features are removed. Rather than solely focusing on the classification task, we presented applications of our methodology to the area of Computational Social Science and confirmed stereotypes such as that @ladygaga (also an LGBT activist) is more popular among @TheDemocrats followers than among @GOP followers. In the domain of marketing we gave evidence that celebrity endorsement is reflected in co-following and we demonstrated how our methodology can be used to reveal the audience similarities between Apple and Puma and, less obviously, between Nike and Coca-Cola. To the best of our knowledge, this is the first systematic study that shows how co-following on Twitter can be used for a variety of applications. Our main focus in this paper was to introduce the concept of (second order) co-following and examine how it works for a wide range of settings, rather than the algorithm itself. In future, we would like to focus more on the algorithmic perspective and extend our work by looking deeper into aspects such as comparing co-following to, e.g., methods which use the tweet content and user profile or compare the plots generated by MDS with other community detection algorithms.

[6] We removed 4 accounts from the initial list that corresponded to media/producers rather than musicians or bands.

[7] `http://en.wikipedia.org/wiki/Avis_Rent_a_Car_System`, accessed on Jan 20, 2014.

7. REFERENCES

[1] Republicans, democrats split on olive garden, other food issues, February 2013. http://newsfeed.time.com/2013/02/27/republicans-democrats-split-on-olive-garden-other-food-issues/.

[2] J. An, M. Cha, P. K. Gummadi, and J. Crowcroft. Media landscape in twitter: A world of new conventions and political diversity. In *ICWSM*, 2011.

[3] P. Barbera. Birds of the same feather tweet together: Bayesian ideal point estimation using twitter data. In *PolNet*, 2013. https://files.nyu.edu/pba220/public/barbera_birds_july2013.pdf.

[4] J. Bernstein, R. E. Bromley, and K. T. Meyer. Republicans and golf, democrats and outkast: Or, party political culture from the top down. *The Forum*, 4(3), 2006.

[5] J. Blakley. The zogby/lear center survey on politics and entertainment, 2008. http://www.learcenter.org/html/projects/?cm=zogby/08.

[6] C. Cioffi-Revilla. Computational social science. *Wiley Interdisciplinary Reviews: Computational Statistics*, 2(3):259–271, 2010.

[7] R. Cohen and D. Ruths. Classifying political orientation on twitter: It's not easy! In *ICWSM*, 2013.

[8] M. Conover, B. Gonçalves, J. Ratkiewicz, A. Flammini, and F. Menczer. Predicting the political alignment of twitter users. In *SocialCom/PASSAT*, pages 192–199, 2011.

[9] J. Davis and M. Goadrich. The relationship between precision-recall and roc curves. In *ICML*, pages 233–240, 2006.

[10] P. Doyle. Favorite beers of republicans and democrats, October 2012. http://www.bostonmagazine.com/news/blog/2012/10/01/favorite-beers-republicans-democrats/.

[11] S. Fortunato. Community detection in graphs. *Physics Reports*, 486(3-5):75–174, 2010.

[12] S. A. Golder and S. Yardi. Structural predictors of tie formation in twitter: Transitivity and mutuality. In *SocialCom/PASSAT*, pages 88–95, 2010.

[13] C. Hutto, S. Yardi, and E. Gilbert. A longitudinal study of follow predictors on twitter. In *CHI*, pages 821–830, 2013.

[14] T. Kahute. Puma brand communication analysis, 2006. http://www.scribd.com/doc/35130565/Puma-Brand-Analysis.

[15] H. Kwak, C. Lee, H. Park, and S. Moon. What is twitter, a social network or a news media? In *WWW*, pages 591–600, 2010.

[16] D. Lazer, A. S. Pentland, L. Adamic, S. Aral, A. L. Barabasi, D. Brewer, N. Christakis, N. Contractor, J. Fowler, M. Gutmann, et al. Life in the network: the coming age of computational social science. *Science*, 323(5915):721, 2009.

[17] M. E. J. Newman. Modularity and community structure in networks. *PNAS*, 103(23):8577–8582, 2006.

[18] Y. Pan, D.-H. Li, J.-G. Liu, and J.-Z. Liang. Detecting community structure in complex networks via node similarity. *Physica A: Statistical Mechanics and its Applications*, 389(14):2849–2857, 2010.

[19] M. Pennacchiotti and A.-M. Popescu. Democrats, republicans and starbucks afficionados: user classification in twitter. In *KDD*, pages 430–438, 2011.

[20] D. D. Posta, Y. Shi, and M. Macy. Why do liberals drink latte?, 2013. under review.

[21] D. Rao, D. Yarowsky, A. Shreevats, and M. Gupta. Classifying latent user attributes in twitter. In *SMUC*, pages 37–44, 2010.

[22] M. Shannon and W. Feltus. What your beer says about your politics, September 2012. http://www.nationaljournal.com/blogs/hotlineoncall/2012/09/what-your-beer-says-about-your-politics-27.

[23] L. Vaughan. Visualizing linguistic and cultural differences using web co-link data. *Journal of the American Society for Information Science and Technology*, 57(9):1178–1193, 2006.

[24] I. Weber and V. R. K. Garimella. Using co-following for personalized out-of-context twitter friend recommendation. 2014.

[25] D. Yin, L. Hong, and B. D. Davison. Structural link analysis and prediction in microblogs. In *CIKM*, pages 1163–1168, 2011.

[26] F. A. Zamal, W. Liu, and D. Ruths. Homophily and latent attribute inference: Inferring latent attributes of twitter users from neighbors. In *ICWSM*, 2012.