

Estimating Tie Strength in Follower Networks to Measure Brand Perceptions

Tung Nguyen, Li Zhang, and Aron Culotta
Department of Computer Science
Illinois Institute of Technology
Chicago, IL 60616

Abstract—As public entities like brands and politicians increasingly rely on social media to engage their constituents, analyzing who follows them can reveal information about how they are perceived. Whereas most prior work considers following networks as unweighted directed graphs, in this paper we use a tie strength model to place weights on follow links to estimate the strength of relationship between users. We use conversational signals (retweets, mentions) as a proxy class label for a binary classification problem, using social and linguistic features to estimate tie strength. We then apply this approach to a case study estimating how brands are perceived with respect to certain issues (e.g., how environmentally friendly is Patagonia perceived to be?). We compute weighted follower overlap scores to measure the similarity between brands and exemplar accounts (e.g., environmental non-profits), finding that the tie strength scores can provide more nuanced estimates of consumer perception.

Index Terms—tie strength, link prediction, public perception

I. INTRODUCTION

Public entities, such as brands, celebrities, and politicians, increasingly use social media to engage with their constituents. Analyzing who follows these entities on sites such as Twitter can reveal information about how they are perceived. For example, two accounts that have many followers in common may be similar in some way [1, 2, 3].

Most analyses of follower networks treat links as binary variables — user A either follows user B or she does not. However, there are many reasons one may follow another user, some indicating a closer relationship than others. This link heterogeneity is further complicated by the growing presence of bots and other inorganic accounts, which can introduce noise into follower analysis.

In this paper, we propose a supervised learning approach to annotate each directed edge in a follower network with a *tie strength* weight. Whereas most prior work in tie strength estimation incorporates user interactions (e.g., retweets, mentions) as *input features* of a tie strength model [4, 5, 6], in our setting these types of interactions are sparse — e.g., most followers

of a brand do not retweet or mention that brand. Instead, we treat user interactions as a weak source of supervision. We formulate a binary classification task to predict whether follower A will retweet or mention another user B , then use the confidence of this classifier on held-out examples as an estimate of tie strength. On a dataset of 85k follower links, we find a random forest achieves AUC scores between .8-.88 on this prediction task.

Additionally, we apply this tie strength method to a case study of estimating the consumer perception of a brand with respect to a given attribute — e.g., which clothing brand is perceived as more eco-friendly, Patagonia or Timberland? Prior work [2] estimated this by analyzing the followers overlap between a brand’s Twitter account and a selected set of *exemplar* accounts representing the attribute of interest (e.g., @Greenpeace and @EPA as exemplars of environmental friendliness). In this paper, we show that including the predicted tie strength weights into the overlap calculation can produce more accurate perception estimates over a range of brand sectors (automobiles, apparel, food and beverage) and attributes (eco-friendliness, nutrition, and luxury).

In the following sections, we first summarize related work regarding this topic (§II) and how our new approach enhances previous findings. §III describes our proposed methodology and §IV describes the social network data we collect for evaluating the approach. We present the experiments and results in §V, followed by a discussion of the implications of the results (§VI).

II. RELATED WORK

Tie strength. Granovetter [7] introduced the concept of *tie strength*, originally to distinguish between strong links (close friends) and weak ties (acquaintances) in offline social networks. Since the introduction of online social networks, a number of methods have been proposed to estimate tie strength between users [4, 8, 5, 9, 10, 6, 11]. These methods construct classification or regression models to predict tie strength based on features of the user pair. Typically, labels are obtained by surveying users about the nature of their relationships, or using existing sources of supervision, such as Facebook’s “Top Friends” list. A critical distinction with the current work is that prior work focuses on relationships that already contain evidences of social interactions (i.e, having direct communication data as one of the features input), and aim

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ASONAM '19, August 27-30, 2019, Vancouver, Canada

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6868-1/19/08?/\$15.00

<https://dx.doi.org/10.1145/3341161.3343675>

to distinguish between friends and acquaintances. As a result, the most predictive features are typically the quantity and type of social interactions between a pair of users. In contrast, our goal in this work is to estimate the tie strength between two users who may not have a history of interaction; i.e., most users who follow a brand do not have direct communications with that brand. In order to measure the tie strength of these asymmetric relationships, our model instead aims to predict the presence of interactions between users, using content and social features. The small number of links that exhibit direct communications serve as a proxy source of supervision for the model. The key assumption is that two users who are more likely to have direct interactions should have a higher tie strength.

Link Prediction. Whereas tie strength models estimate weights for existing edges, link prediction models instead predict the creation of edges that do not yet exist [12, 13, 14, 15]. In our case study of brand perceptions, the similarity computation between a brand and an exemplar can be viewed as a type of link prediction problem. E.g., to determine how eco-friendly @Patagonia is perceived to be, we can estimate the likelihood of an edge connecting @Patagonia to an eco-friendly exemplar like @Greenpeace. Below, we propose a weighted variant of the common Jaccard similarity measure to compute the similarity of two sets of followers. Zhu and Xia [16] propose a weighted neighbor overlap function that has some similarities to our weighted Jaccard measure; however, the edge weights are produced using only graph properties, and the method is not applied to online social networks.

Brand Perception. Tracking public perceptions is a key issue in marketing, politics, and public health. Typically, this is done with surveys and focus groups; however, surveys can be costly and time-consuming to conduct, making them particularly ill-suited to rapidly changing environments. Furthermore, survey response rates have fallen precipitously in recent years [17], leading researchers to seek non-traditional methods to gauge public perception. An emerging alternative to traditional surveys is online social network analysis. If successful, such an approach would provide real-time insight into evolving perceptions, enabling fine-grained stratification of results across population groups.

Most existing algorithmic measures of public perception rely on text-based sentiment analysis [18, 19, 20, 21], which is critically limited by the scarcity of explicit expressions of opinion and the difficulty of accurately parsing linguistic nuances. Recently, however, new techniques have been proposed that measure public perception of an entity by analyzing the social network structure around it [2, 22]. For example, Culotta and Cutler [2] found that if many of a brand’s Twitter followers also follow pro-environmental accounts, then it is more likely that the brand is perceived as eco-friendly by consumers, as validated using survey data.

However, this prior work assumes all the links between brands and their followers have the same weights, ignoring the distinction between strong and weak ties. In this paper, we develop data-driven tie strength scores, assign them as

edge weights in the follower graph, and use those to compute new measures of perception. We find that this more nuanced measure often produces more accurate measures of consumer perception.

III. METHODS

A. Tie strength estimation

Given a following relationship $u_1 \rightarrow u_2$, our goal is to estimate a tie strength score $T(u_1, u_2) \in [0, 1]$, where larger scores indicate a closer affinity between u_1 and u_2 . Whereas prior work in tie strength estimation assumes access to labeled data (e.g., surveys indicating whether u_1 and u_2 are close friends or acquaintances) [23, 4, 8], such data is unavailable at the scale required for our application. Furthermore, in our case u_2 is typically a brand or other organizational account, and u_1 is an “ordinary” user. Thus, the friend/acquaintance distinction in prior work is not quite appropriate for this asymmetric relationship.

Instead, we propose using a surrogate source of supervision based on the social interactions between two users. Specifically, we formulate a binary classification task where we label a link $u_1 \rightarrow u_2$ as a positive example if u_1 retweets or mentions u_2 , or if u_1 uses a hashtag that is the username of u_2 . Thus, we view these social interactions as indicators of a stronger relationship between u_1 and u_2 . For example, if one user follows @Greenpeace and retweets @Greenpeace, this is considered a stronger relationship than for another user who follows @Greenpeace but never retweets them.

The justification for this model choice comes from prior work [4, 8, 24, 6] indicating that one of the most important factors in tie strength is interaction data, or the communication between two users. In our case, this type of data is sparse since, in contrast to friends, users will rarely have a constant communication with a brand. On Twitter, there are three main methods of communicating with public profiles: retweet, mention and hashtag. While prior work suggests that these three forms of interaction can have slightly different implications for the relationship [24], given their overall sparsity in this data, we consider the presence of any of the three as an indicator of a strong relationship.

Formally, our approach is to fit a binary classifier to the labeled data we have, then apply the classifier to each edge in the following network to assign tie strength scores. Given a set of accounts $U = \{u_1 \dots u_n\}$, let F_{u_i} indicate the followers of user u_i , and let $u_i^j \in F_{u_i}$ indicate the j th follower of user u_i . The binary classification examples are then tuples of the form $\{(u_1, u_1^1, y_1^1) \dots (u_i, u_i^j, y_i^j) \dots (u_n, u_n^k, y_n^k)\}$ where $y_i^j = 1$ if user u_i^j interacts with user u_i through a retweet, mention, or hashtag; otherwise $y_i^j = 0$.

For each tuple (u_i, u_i^j, y_i^j) , we compute a number of features to be used in a binary classifier to predict the value of y_i^j . These features are designed to indicate linguistic and social similarities between the brand/exemplar u_i and the follower u_i^j . (Note that any tweets that are used to derive the class label are removed from the input features for the purpose of

classification.) Below, we summarize the features used for this task.

- 1) **Tweets similarity:** For each follower u_i^j and organizational account u_i , we collect their most recent 200 tweets (including retweets and replies) from their timelines using the Twitter API. We concatenate these tweets into a single bag-of-words vector, using tf-idf transformation, then compute the cosine similarity between the two term vectors. We compute two features: one using unigrams and one using bigrams (**tweets similarity 1-gram** and **tweets similarity 2-gram**).
- 2) **Tweets entities similarity:** Results in Macskassy and Michelson [25] suggest that instead of using the entire text of users’ tweets, using only important named entities can result in more predictive features. While there is a long line of research into named entity recognition, we find that the simple heuristic proposed in Macskassy and Michelson [25] works well — after removing stop words and URLs, we identify all capitalized words as named entities. As in the tweets term similarity features, we then represent each user as a bag of entities, and compute the cosine similarity between users using both unigrams and bigrams.
- 3) **Favorite tweets similarity:** We also collect the tweets that each user has “favorited.” We hypothesize that such tweets are more indicative of the tastes and preferences of the user. To capture this, we perform the same calculations as in the **tweets similarity** features above, but only on the tweets that each user has favorited. We similarly compute unigram and bigram versions of this feature (**fav. tweets similarity 1-gram** and **fav. tweets similarity 2-gram**).
- 4) **Favorite tweets entities similarity:** We repeat the **tweets entity similarity** computation using only the favorited tweets from each user (**fav. tweets entities 1-gram** and **fav. tweets entities 2-gram**).
- 5) **Description similarity:** We repeat the **tweets similarity** features instead using the description field of each user (**description entities 1-gram** and **description entities 2-gram**).
- 6) **Number of friends overlap:** We compute the overlap of the accounts that each user follows. Letting $N(u_i)$ be the set of users that u_i follows, we compute the Jaccard similarity $J(N(u_i), N(u_i^j))$.
- 7) **Number of favorite accounts overlap:** Similarly, we compute the overlap of the accounts that each user has favorited (when a user u_i favorites a tweet posted by user u_j , we say that user u_i has favorited user u_j). We again use the Jaccard similarity to compute this feature.
- 8) **Number of follows in same sector ratio:** We divide all the brands into 4 different sectors: Food, Personal Care, Car, and Apparel. The set of exemplars is also separated into 3 different sectors: Eco-friendliness, Luxury, and Nutrition. When u_i is a brand, we compute a feature indicating the fraction of brands in the same sector as u_i that user u_i^j follows. Similarly, when u_i is an exemplar, we compute the fraction of exemplars in the same sector as exemplar u_i

that user u_i^j follows. This feature indicates how connected user u_i^j is to the rest of the accounts in the sector.

Given the training data and the features defined above, we fit a random forest classifier [26] to perform binary classification. To compute the tie strength for each follow link $u_i^j \rightarrow u_i$, we perform 10-fold cross validation over all data, and set the tie strength equal to the posterior probability $P(y = 1 | u_i, u_i^j)$, as computed by the random forest trained on a disjoint training set. We refer to this tie strength score as $T(u_i, u_i^j)$.

B. Case Study: Social Perception Scores

Given these tie strength estimates, we can now construct a weighted follower network containing users, brands, and exemplar accounts. In this section, we propose methods to analyze this data structure to estimate the consumer perception of brands.

For an online organizational account u (e.g., a brand or politician) and a perceptual attribute a (e.g., eco-friendliness), our goal is to compute a **perception score** $s(u, a) \in \mathbb{R}$ that quantifies the public perception of u with respect to attribute a , where higher values indicate a more positive perception. For example, if u_1 and u_2 are the accounts of two car brands, and a indicates environmental friendliness, then $s(u_1, a) > s(u_2, a)$ indicates that brand u_1 is perceived as more environmentally friendly than brand u_2 .

In the next two sections, we first review prior work that computes a perception score from social network information, then proceed to introduce a new approach based on tie strength. The discussion and experiments below will assume Twitter as the data source, although it is possible to extend the approach to other networks that support “follow” and “like” connections.

1) *Jaccard Social Perception Score:* Culotta and Cutler [2] proposed a method of estimating perception scores based on the affinity of a user to a list of *exemplar* accounts. An exemplar is an account that is known to be strongly affiliated with attribute a . For example, Environmental Protection Agency and Sierra Club may be said to exemplify the eco-friendliness attribute. Let $e \in E$ be an exemplar for attribute a ; the approach of Culotta and Cutler [2] assumes that if a user u is followed by users who also follow exemplars in E , then u is likely to be perceived strongly with respect to a . For example, if a large fraction of Toyota’s followers on Twitter also follow the Sierra Club, then it is likely that Toyota is perceived as eco-friendly. In other words, a user is perceived to be eco-friendly if people who value eco-friendliness support that user.¹

The specific formula they use is based on the Jaccard similarity between the followers of a brand and the exemplar accounts. Let F_i be the set of followers of user i , and let

¹The most likely data generation process is that a brand is perceived as eco-friendly, and then eco-friendly consumers decide to follow that brand online (not the reverse).

$J(A, B)$ be the Jaccard score between sets A and B , then the **Jaccard Social Perception Score (Jaccard SPS)** is:

$$s_J(u, a, E) = \frac{1}{\sum_e \frac{1}{|F_e|}} \sum_{e \in E} \frac{1}{|F_e|} J(F_u, F_e) \quad (1)$$

This is the weighted average of the Jaccard similarity between the followers of user u and the followers of each exemplar e . The weights are the inverse of the number of followers for each exemplar; the intuition is that exemplars with few followers have a more focused audience which is likely to be more indicative of the attribute of interest. Experiments also found applying square root to the final score was useful to mitigate the effects of positive skew in the scores.

While this approach makes a number of assumptions about the many possible reasons a user decides to follow another user online, it is supported by a wealth of research indicating the prevalence of homophily in social networks [27, 28, 29]. The results in Culotta and Cutler [2] showed strong correlations between the perception scores and the results of consumer surveys of brands (.5-.8 correlation).

One limitation of this approach is that it assumes each following link has the same strength. However, in practice we know that different users have different tie strengths to the users they follow. For example, some users who follow exemplars like Greenpeace may be strongly involved in environmental activism, while others may only have a passing interest. Similarly, users who follow brands like Toyota may have strong or weak connections to the brand. Using the tie strength estimates derived above, we propose a weighted variant of Jaccard similarity to estimate public perception.

2) *Tie Strength Social Perception Score*: Letting $T(u, v)$ be the tie strength between users u and v , and letting F_u be the set of followers of user u , we first define the **total follower strength** of a brand/exemplar u as

$$T_*(u) = \sum_{v \in F_u} T(u, v) \quad (2)$$

This quantity summarizes both the number of followers of u as well as their tie strength. While this quantity may have value on its own from a brand’s perspective (e.g., to measure the quality of their online followers), here we use it as a normalizing factor when computing follower similarities.

We define the **Tie Strength Social Perception Score (Tie Strength SPS)** as follows:

$$s_T(u, a, E) = \sum_{e \in E} \frac{\sum_{o \in F_u \cap F_e} T(o, u) * T(o, e)}{T_*(u) + T_*(e)} \quad (3)$$

This formula sums together the weighted similarities between each exemplar $e \in E$ with the brand u . It is a weighted analog to the Jaccard similarity in that the numerator considers users who follow both the exemplar e and the brand u ; rather than just computing the size of the intersection, we multiply together the tie strength scores for each user in the intersection. Similarly, the denominator accounts for the overall tie strength for all followers of u and e , analogous to how the Jaccard

similarity uses the size of the union of the two sets. This reduces the bias towards brands with a large number of followers. Thus, for brand u to have a strong relationship with exemplar e , users o who follow both u and e must have strong tie strengths with them.

Finally, we consider an ensemble score that combines the approaches of Equations 1 and 3:

$$s(u, a, E) = \sqrt{c_1 * s_T(u, a, E)} + \sqrt{c_2 * s_J(u, a, E)} \quad (4)$$

where c_1 and c_2 are constants weighting the importance of each perception score, and where we use a square root transformation to mitigate the impact of positive skew in the scores.² We refer to this final score as the **Ensemble Social Perception Score (Ensemble SPS)**.

IV. DATA

In this section, we describe the data we collected to compute tie strength and social perception scores, as well as the survey data used to validate the results.

A. Brands and exemplars data

For a fair comparison, we begin with the same list of brands and exemplars from Culotta and Cutler [2]. These data contain Twitter user ids for 154 brands identified on the website GoodGuide.com, coming from four sectors: Car, Apparel, Food & Beverage and Personal Care.

We consider three perceptual attributes: environmental friendliness, nutrition, and luxury. For each attribute, we use the list of exemplars from prior work, which were discovered by searching for users on Twitter Lists relevant to each attribute. These consist of 69 eco-friendly, 360 nutrition and 95 luxury exemplars.

For each brand and exemplar, we use the list of followers from prior work. However, in order to compute our tie strength scores, we use the Twitter API to collect additional information for each organizational account: its most recent 200 tweets (including retweets and replies), its description, its list of 200 most recent favorite (or liked) tweets and their corresponding authors, and its list of at most 5,000 friends (i.e., accounts that the brand follows).

B. Followers Data

To compute Jaccard SPS, we use all the follower ids collected, roughly 35M total. For Tie Strength SPS, however, we need to also collect additional information on each follower (e.g., profile, recent tweets, likes, etc.). As this is infeasible to collect for all followers due to rate limitations, we instead sample uniformly a smaller set of users to collect additional data to compute tie strength. To do so, for each brand-exemplar pair, we sample 10 overlapping followers for tie strength computation. Additionally, for tie strength experiments, we sample an equal number of positive and negative examples for each brand/exemplar. After filtering for bots and celebrity accounts, this results in 85,888 followers across all attributes, for which

²In the experiments below, we set $c_1 = 10^5$ and $c_2 = 10^9$ for all results; preliminary experiments suggest the method is fairly robust to these values.

	AUC		
	eco	nutrition	luxury
Random Forest Classifier	0.88	0.80	0.85
Logistic Regression	0.69	0.68	0.66
Neural Network	0.79	0.76	0.76

TABLE I: Average AUC of different tie strength estimation models (highest values in bolds).

	Features importance		
	eco	nutrition	luxury
Tweets entities 1-gram	0.18	<i>0.12</i>	0.22
Tweets entities 2-gram	0.04	0.05	0.05
Fav. tweets entities 1-gram	<i>0.16</i>	0.08	<i>0.20</i>
Tweets similarity 1-gram	0.07	0.09	0.09
Tweets similarity 2-gram	0.10	0.11	0.08
Friends overlap	0.07	0.10	0.09
Favorites overlap	0.05	0.14	0.09
Fav. tweets similarity 1-gram	0.14	0.08	0.09
Fav. tweets similarity 2-gram	0.11	0.09	0.07

TABLE II: Feature importances for random forest models (top feature in bold, second highest in italics).

we collect their most recent 200 tweets, 200 favorited tweets and their authors, and up to 5,000 friends. These users are used to train and evaluate the tie strength model and to compute Tie Strength SPS.

C. Survey Data

To validate our perceptions scores, we use the brand perception survey results from Culotta and Cutler [2]. The survey was conducted with 500 participants recruited through Amazon Mechanical Turk (AMT). These participants were all located in the United States with high records on AMT (having done at least one hundred assignments with a 95% acceptance rate) to strengthen the quality of the results. The participants' task was to rate each brand on a scale of one to five indicating how strongly they view the brand with respect to each perceptual attribute. Each person rated a random sample of between 40 to 70 brands of the same sector or attribute. After filtering out invalid responses, the survey contains results for 340 participants. To validate the automatically computed perception scores, we compute the Spearman correlation with the survey scores, as the relative brand ranking is more useful in this domain than the absolute score. (Results using Pearson correlation are similar.)

V. EXPERIMENTS AND RESULTS

In this section, we first describe results validating the tie strength scores, then compare the impact this has on the final perception scores.

A. Tie strength estimation

For each attribute, we perform cross-validation to tune hyper-parameters and perform feature selection separately for

	Details	
	Brand	
User 1	Brand	Naked Juice / tie strength=0.99
	Sample tweet	Community Winners #SFGreenFest #SF go for a hike&learn about the environment! @DotGreenTLD
	Description	Green. Health. Art. Inspiration. Environment. Internet. Fun!!
User 2	Brand	Lamborghini / tie strength=0.99
	Sample tweet	RT @DoctaM3: What events should we do this year? Lots of BIG things planned for the Lamborghini Huracan Performante aka... https://t.co/mvT
	Description	Fast cars, family and business. Co-Founder of @ShipsALot. Founded @JerkyXP. Sometimes @JackTheCEO let's me pet him.
User 3	Brand	Nike / tie strength=.93
	Sample tweet	Blazers need a G-League team to develop their young players, instead of having them ride the bench
	Description	Former News & Sports Reporter, Korean, Fitness, MMA, EDM

TABLE III: A sample of users and the brand they follow with high estimated tie strength.

each attribute. The Eco-friendliness model uses at most 20 trees in the forest, a minimum of 5 nodes per leaf, and the maximum tree depth is 50. On the other hand, Nutrition and Luxury models use at most 42 trees in the forest, a minimum of 5 nodes per leaf with maximum depth of 42.

Our full feature set contains 13 total features (§III-A). We conduct an exhaustive enumeration of all feature combinations to assess the impact of each. Results show that none of the best models will need to utilize all features. For nutrition, 11 features lead to the highest classification accuracy; whereas only 10 features are required for luxury and eco-friendliness. However, the optimal set of utilized features are different across datasets. Table II shows the list of features that appear in all 3 models with their corresponding feature importance scores provided by the Gini impurity measure for random forests (as implemented in scikit-learn [30]). These results indicate that the most predictive features include tweets entities, favorite tweets entities, and tweets' content similarity. It is notable that the favorite and entity variants add predictive power over using the raw tweet content alone, reaffirming that these more specialized features better reflect a user's tastes and preferences.

For the classifier, we choose to use Random Forest Classifier. Preliminary experiments show that Logistic Regression and Neural Networks yield lower accuracies compared to our technique. Results of these experiments are listed in Table I, which shown the AUCs averaged over the 10 cross-validation folds for each attribute. The values indicates that our chosen method provides a highly predictive model, with the lowest AUC being 0.8 on nutrition. The other two AUCs are 0.88 on eco-friendliness and 0.85 on luxury, suggesting a high

Attribute	Brand	Tie strength (Rank)
Eco-friendliness	Nature Valley	42.88 (5)
	Sprite	15.63 (39)
	Kia	216.17 (5)
	Cadillac	147.28 (25)
Nutrition	Naked Juice	23.87 (4)
	SunChips	14.37 (17)
Luxury	Ann Taylor	101.48 (5)
	Hollister Co.	79.07 (15)
	Diesel	53.87 (8)
	Aeropostale	33.44 (22)

TABLE IV: Total tie strength scores (and corresponding survey ranking) between pairs of brands that have similar number of overlapping followers with the exemplars.

accuracy in predicting positive instances by our model.

Additionally, Table III shows an overview of the users with the highest tie strength score to a specific brand. Sampled tweets show a high overlap between the user’s Twitter profile and the brand’s social media contents. For example, user 2 shows interest in cars and technology both in his description and his tweets/retweets contents. Hence, it is reasonable to conclude that this user will have a high tie strength to Lamborghini, a luxury car brand that he follows.

We also investigate how the predicted tie strength scores replicate and add nuances to the exploration of brand-followers relationships. In the experiment, we group brands that had similar number of overlapping followers with exemplars, but different total following strength (Equation 2). For example, Nature Valley and Sprite have approximately the same number of overlapping followers with eco-friendliness exemplars (94 and 91). However, the total edge scores of the two brands differ substantially; Nature Valley has a score of 42.88 versus 15.63 for Sprite. This is in line with the survey results, in which consumers rate Nature Valley as more eco-friendly than Sprite (rank 5 vs 39). Those results suggest that the tie strength scores can help provide a more fine-grained affinity between a brand and an exemplar, particularly when the number of overlapping users is small. Table IV gives more examples of these brand pairs, which indicates that tie strength scores are important across all perceptual attributes.

Finally, we conduct an error analysis to understand the shortcomings of the tie strength model. We find that for many of the false positives the user and the brand/exemplar have similar characteristics, yet the user has never interacted with the brand/exemplar. E.g., both may post frequently about environmental issues, but have never interacted before. This is in part a trade-off of using the proxy class labels to train the classifier — there will be some noise in the class labels when a user in fact has a strong tie to the account, but there do not exist social interactions in our data to support it.

Attribute	Ensemble SPS	Tie strength SPS	Jaccard SPS (baseline)
Car/Eco	0.670	0.231	<i>0.601</i>
Apparel/Eco	0.458	0.238	<i>0.419</i>
Food&Bev./Eco	<i>0.735</i>	0.071	0.750
Personal Care/Eco	<i>0.525</i>	0.583	0.511
Food&Bev./Nutrition	0.843	0.522	<i>0.839</i>
Apparel/Luxury	0.713	0.500	<i>0.679</i>
Car/Luxury	0.541	0.416	<i>0.522</i>

TABLE V: Spearman correlation for competing algorithms.(top score in bold, second highest in italics)

B. Public perception estimation

After using the tie strength classifier to assign weights to each edge in the follower network, we proceed to compute and validate the final social perception score for each of our 154 brands. The goal of our experiments is to determine the impact of the tie strength scores on the final correlation with the survey results.

Table V presents the Spearman correlations between multiple SPS methods and the survey scores. From left to right, the Ensemble SPS (Equation 4) represents our method’s rating system, the Tie strength SPS (Equation 3) is the score calculated solely from our tie strength estimation method, and the Jaccard SPS is the baseline model which utilize Jaccard similarity (Equation 1, with square root transformation). The results show that the Ensemble SPS provides higher correlation with the survey results than the baseline across 6 of 7 datasets. The largest increase comes from Car/Eco-friendliness pair with a 6.9% increase; while the only section Ensemble SPS performs worse is Food & Beverages / Eco-friendliness (a 1.5% decrease from 0.75 to 0.735). The lower overall scores of Tie Strength SPS is most likely due to the small sample size of followers it uses; however, it is notable that combining it with Jaccard SPS produces the most accurate estimates. The Spearman correlation between Ensemble SPS and the survey is visualized in Figure 1.

For a subjective evaluation, we present the top and bottom ranked brands for each attribute in Table VI. These provide anecdotal evidences of the quality of the scoring method — e.g., brands like “Organic Valley” and “Nature’s Path“ are more closely aligned with health and nutrition than brands like “Redbull” or “Starburst”.

Finally, Figures 2 and 3 show the tie strength distribution of brands with the highest change in ranking when using Ensemble SPS instead of Jaccard SPS. From the figure, we can see that brands with an increased ranking have positively skewed distributions. On the other hand, brands with a lower ranking have negatively skewed distributions. Moreover, it is noticeable from the figures that brands with higher rank increase tend to have more tie strength scores link to them.

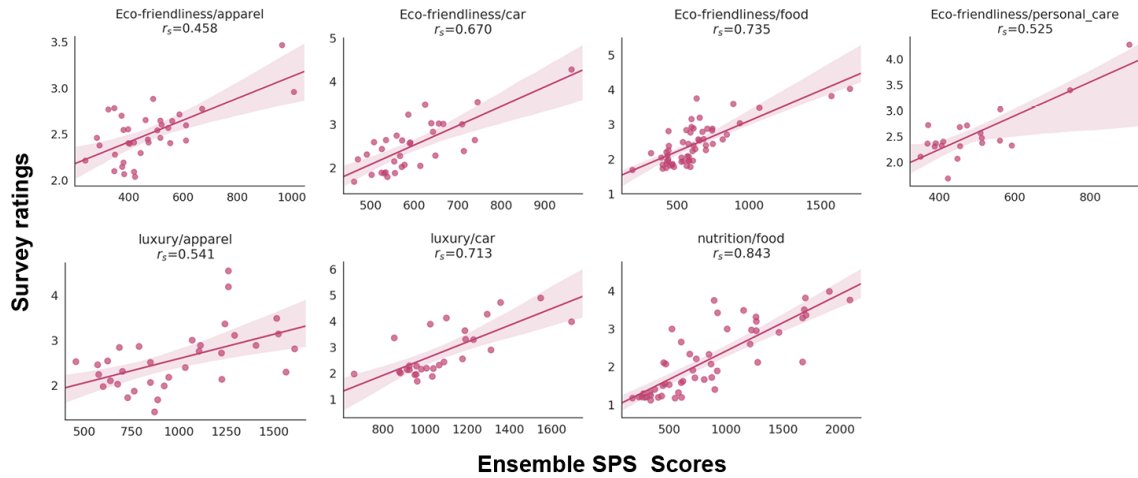


Fig. 1: Spearman correlation between Ensemble SPS scores and survey ratings on different sector/attribute pairs.

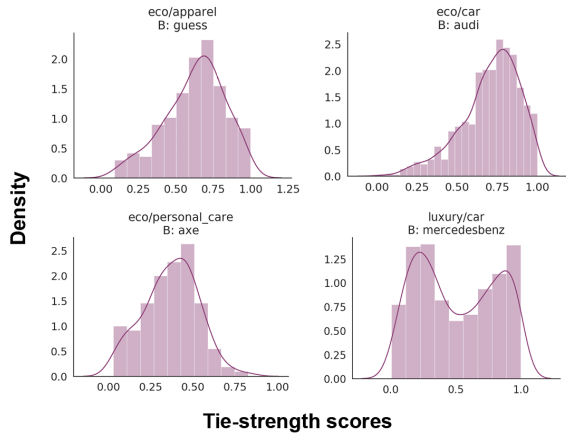


Fig. 2: Tie strengths distribution for brands with large rank increases (compared to Jaccard SPS).

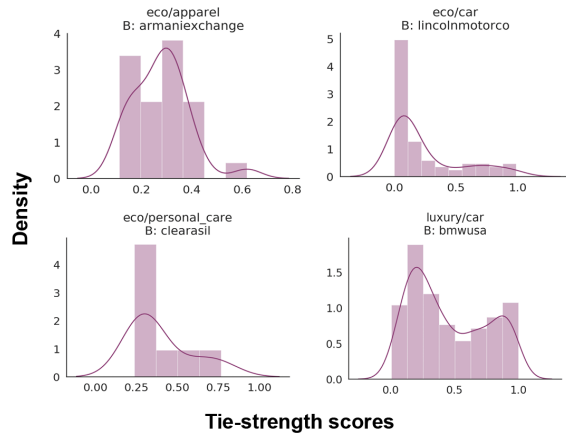


Fig. 3: Tie strengths distribution for brands with large rank decreases (compared to Jaccard SPS).

Attribute	Top brands	Bottom brands
Car/Eco	SmartUSA, Ford, Toyota	Pontiac, Lamborghini
Apparel/Eco	The North Face, Timberland	Aeropostale, Hollister Co.
Food&Bev/Eco	Organic Valley, Nature's Path	Monster Energy, Cadbury, Sprite
Per.Care/Eco	Burt's Bees, Aveda, Suave	Proactiv, VO5, L'Oreal
Food/Nutrition	Organic Valley, Nature's Path	Monster Energy, Redbull, Starburst
Apparel/Luxury	jCrew, Ann Taylor	Nike, iZod, Aeropostale
Car/Luxury	BMW, Rolls Royce, Aston Martin	Pontiac, Smart USA, Lincoln Motor

TABLE VI: Top and bottom brands ranked by the Ensemble SPS for each sector/attribute pair

These results can be explained by the amount of social engagements of a brand, which lead to changes in the final social perception calculation.

VI. CONCLUSIONS AND FUTURE WORK

A primary contribution of this paper is the use of a random forest classifier to predict tie strengths between users and brands/exemplars. Experiments on 85k follower links suggest our approach is effective at estimating the tie strength between Twitter users, with AUC values between .80-.88. We believe this classification model contributes a new, fine-grained approach to the study of online user behavior. Brands and organization can also use this model to gain more insights on their followers and targeted customers.

We then formulated a rating system for the relationship between a brand and an exemplar of the targeted attribute, called the Ensemble Social Perception Score (Ensemble SPS). The Ensemble SPS utilizes both the Jaccard Social Perception

Scores formula presented in Culotta and Cutler [2] and the Tie strength Social Perception Scores (Tie strength SPS) calculated by the tie strength scores of each brand/exemplar. Comparisons with consumer surveys indicate that the tie strength scores provide a more nuanced measure of social similarity.

We also acknowledge some limitations of our approach. One of the limitations stems from the amount of fine-grained users data we have. Due to the Twitter API rate limit, we are limited in the number of links for which we can apply the tie strength classifier. However, we found that even with a small sample of users, we can compute more reliable brand perception scores.

Another limitation is the sensitivity of the approach to accounts with few overlapping followers. In extreme cases, there may be only a few users that follow both a brand and an exemplar. This can lead to high variance estimates of the affinity between the brand and the exemplar; in future work, we will consider additional measures to smooth these values, by considering indirect connections between users.

Future work should consider how such models can be applied to different online social networks. Facebook, Reddit, and Yelp all have different notions of user interactions, and so the method may need to be adapted to the types of interactions available.

ACKNOWLEDGMENTS

This research was funded in part by the National Science Foundation under awards #IIS-1526674 and #IIS-1618244.

REFERENCES

- [1] J. An, M. Cha, K. Gummadi, and J. Crowcroft, "Media landscape in twitter: A world of new conventions and political diversity," in *Fifth International AAI Conference on Weblogs and Social Media*, 2011.
- [2] A. Culotta and J. Cutler, "Mining brand perceptions from Twitter social networks," *Marketing Science*, vol. 35, no. 3, pp. 343–362, 2016.
- [3] J. Peng, A. Agarwal, K. Hosanagar, and R. Iyengar, "Network overlap and content sharing on social media platforms," *Journal of Marketing Research*, vol. 55, no. 4, pp. 571–585, 2018.
- [4] E. Gilbert and K. Karahalios, "Predicting tie strength with social media," in *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 2009, pp. 211–220.
- [5] R. Xiang, J. Neville, and M. Rogati, "Modeling relationship strength in online social networks," in *Proceedings of the 19th international conference on World wide web*. ACM, 2010, pp. 981–990.
- [6] J. J. Jones, J. E. Settle, R. M. Bond, C. J. Fariss, C. Marlow, and J. H. Fowler, "Inferring tie strength from online directed behavior," *PloS one*, vol. 8, no. 1, p. e52168, 2013.
- [7] M. S. Granovetter, "The strength of weak ties," *American journal of sociology*, pp. 1360–1380, 1973.
- [8] I. Kahanda and J. Neville, "Using transactional information to predict link strength in online social networks," in *Third International AAI Conference on Weblogs and Social Media*, 2009.
- [9] E.-A. Baatarjav, A. Amin, R. Dantu, and N. K. Gupta, "Are you my friend?[twitter response estimator]," in *2010 7th IEEE Consumer Communications and Networking Conference*. IEEE, 2010, pp. 1–5.
- [10] E. Gilbert, "Predicting tie strength in a new medium," in *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*. ACM, 2012, pp. 1047–1056.
- [11] H. Huang, Y. Dong, J. Tang, H. Yang, N. V. Chawla, and X. Fu, "Will triadic closure strengthen ties in social networks?" *ACM Transactions on Knowledge Discovery from Data (TKDD)*, vol. 12, no. 3, p. 30, 2018.
- [12] D. Liben-Nowell and J. Kleinberg, "The link-prediction problem for social networks," *Journal of the American society for information science and technology*, vol. 58, no. 7, pp. 1019–1031, 2007.
- [13] A. Clauset, C. Moore, and M. E. Newman, "Hierarchical structure and the prediction of missing links in networks," *Nature*, vol. 453, no. 7191, p. 98, 2008.
- [14] L. Lü and T. Zhou, "Link prediction in weighted networks: The role of weak ties," *EPL (Europhysics Letters)*, vol. 89, no. 1, p. 18001, 2010.
- [15] M. Araújo, P. Ribeiro, and C. Faloutsos, "Tensocast: forecasting time-evolving networks with contextual information," in *Proceedings of the 27th International Joint Conference on Artificial Intelligence*. AAAI Press, 2018, pp. 5199–5203.
- [16] B. Zhu and Y. Xia, "Link prediction in weighted networks: A weighted mutual information model," *PloS one*, vol. 11, no. 2, p. e0148265, 2016.
- [17] A. Kohut, "Assessing the representativeness of public opinion surveys," *Pew Research Center*, 2012.
- [18] B. Pang and L. Lee, "Opinion mining and sentiment analysis," *Foundations and trends in information retrieval*, vol. 2, no. 1-2, pp. 1–135, 2008.
- [19] P. S. Fader and R. S. Winer, "Introduction to the special issue on the emergence and impact of user-generated content," *Marketing Science*, vol. 31, no. 3, pp. 369–371, 2012.
- [20] P. Saleiro, S. Amir, M. Silva, and C. Soares, "Popmine: Tracking political opinion on the web," in *IEEE International Conference on CIT/IIUC/DASC/PICOM*. IEEE, 2015, pp. 1521–1526.
- [21] A. Xu, H. Liu, L. Gou, R. Akkiraju, J. Mahmud, V. Sinha, Y. Hu, and M. Qiao, "Predicting perceived brand personality with social media," in *ICWSM*, 2016, pp. 436–445.
- [22] A. Plumeyer, P. Kottemann, D. Böger, and R. Decker, "Measuring brand image: a systematic review, practical guidance, and future research directions," *Review of Managerial Science*, pp. 1–39, 2017.
- [23] F. Liberatore and L. Quijano-Sanchez, "What do we really need to compute the tie strength? an empirical study applied to social networks," *Computer Communications*, vol. 110, pp. 59–74, 09 2017.
- [24] P. A. Grabowicz, J. J. Ramasco, E. Moro, J. M. Pujol, and V. M. Eguiluz, "Social features of online networks: The strength of intermediary ties in online social media," *PLoS ONE*, vol. 7, no. 1, p. e29358, Jan. 2012.
- [25] S. A. Macskassy and M. Michelson, "Why do people retweet? anti-homophily wins the day!" in *ICWSM*, 2011, pp. 209–216.
- [26] L. Breiman, "Random forests," *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [27] J. E. Lydon, D. W. Jamieson, and M. P. Zanna, "Interpersonal similarity and the social and intellectual dimensions of first impressions," *Social Cognition*, vol. 6, no. 4, pp. 269–86, 1988.
- [28] M. McPherson, L. Smith-Lovin, and J. M. Cook, "Birds of a feather: Homophily in social networks," *Annual Review of Sociology*, vol. 27, no. 1, pp. 415–444, 2001.
- [29] H. G. Pereira, M. de Fátima Salgueiro, and I. Mateus, "Say yes to Facebook and get your customers involved! Relationships in a world of social networks," *Business Horizons*, vol. 57, no. 6, pp. 695–702, 2014.
- [30] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg *et al.*, "Scikit-learn: Machine learning in python," *Journal of machine learning research*, vol. 12, no. Oct, pp. 2825–2830, 2011.