Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI-16)

Predicting Gaming Related Properties from Twitter Accounts

Maria Ivanova Gorinova
University of Cambridge
Cambridge, UK

Yoad Lewenberg
The Hebrew University of Jerusalem
Jerusalem, Israel

Yoram Bachrach
Microsoft Research
Cambridge, UK

Alfredo Kalaitzis
Microsoft
London, UK

Michael Fagan
Microsoft
London, UK

Dean Carignan
Microsoft
Redmond, US

Nitin Gautam
Microsoft
Redmond, US

Abstract

We demonstrate a system for predicting gaming related properties from Twitter accounts. Our system predicts various traits of users based on the tweets publicly available in their profiles. Such inferred traits include degrees of tech-savviness and knowledge on computer games, actual gaming performance, preferred platform, degree of originality, humor and influence on others. Our system is based on machine learning models trained on crowd-sourced data. It allows people to select Twitter accounts of their fellow gamers, examine the trait predictions made by our system, and the main drivers of these predictions. We present empirical results on the performance of our system based on its accuracy on our crowd-sourced dataset.

Introduction

Social media plays an important role in our lives, and services such as Twitter, Facebook and Google+ are used regularly by over a billion users. Recent research has uncovered many ways in which online information, including social network data, can be used to predict personal traits of users (Golbeck et al. 2011; Bachrach et al. 2012; Kosinski et al. 2012; 2014; Kosinski, Stillwell, and Graepel 2013; Schwartz et al. 2013; Bachrach 2015),

Such information can be used to gain insight regarding users (Bachrach et al. 2014c; Lewenberg, Bachrach, and Volkova 2015), or for commercial applications such personalized search (Ustinovsky and Serdyukov 2013), targeted advertising (Yang et al. 2006) or improving the quality of collaborative filtering based recommender systems.[1]

Such earlier work focuses on the general population, whereas our focus is on the specific target group of people playing computer games. Gamers are predominantly active in social media, and use distinct online communication styles and language. We focus on the following *perceived* traits of gamers, assumed to affect their standing in the gaming community: tech-savviness, degree of knowledge on computer games, and gaming skill in various genres. We also examine variables such as their life-stage, degree of originality, and level of influence on their peers.

Given a Twitter handle, our system predicts the user's traits. These predictions are the result of applying machine learning classifiers, which examine the textual tweets made by the target user.

[1]Standard recommender systems only use information on consumed items and rely on fingerprinting or dimensionality reduction techniques (Koren, Bell, and Volinsky 2009; Bachrach et al. 2014b; Bachrach and Porat 2015), but can be adapted to incorporate more detailed user profiling (Clements et al. 2011; Bachrach et al. 2014a).

Methodology

We build a model for each of the gaming-related traits mentioned above and train it on a dataset of 2,000 Twitter accounts, annotated by workers on Amazon's Mechanical Turk. Each worker was asked to examine several of those Twitter accounts and form an opinion regarding the traits of the account owners. Each Twitter profile was annotated multiple times, on average by 3 raters. Further, we examine the amount of time that these target users have spent playing Xbox games, and their actual achievement scores.

The textual data of the users in the training dataset is preprocessed by reducing all words to their root form, using a Porter Stemmer. The stemmed text is then used to extract a vocabulary, which consists of those words and hashtag words (those prefixed by '#') that are present in at least 3 user profiles and at most 80% of all user profiles.

We combine *lexical* and *stylistic* features to train the models and classify new users. The *lexical* part of a feature vector is obtained by applying TF-IDF weighting to the users' tweets, with respect to the extracted vocabulary. The *stylistic* features include occurrences of elongated words, fully capitalized words, consecutive punctuation marks, hashtags, as well as the percentage of the tweets that were retweets or replies, and the number of URLs that the user shared.

Feature extraction and classification

We use linear regression for quantitative traits and logistic regression for classification.

Our system inputs a Twitter username, scrapes the recent tweets of the user via the Twitter API and automatically translates them through the Microsoft Translator API where appropriate.

	Gender	Age	Life stage	Tech
ICC	0.92	0.71	0.69	0.55
RMSE	0.52	0.41	0.42	0.39
Accuracy	73%	83%	83%	85%
	Knowledge	**Trustworthy**	**Quality**	**Funny**
ICC	0.76	0.37	0.38	0.33
RMSE	0.26	0.45	0.46	0.47
Accuracy	93%	79%	79%	77%
	Originality	**Influencial**	**Xbox fan**	**PS fan**
ICC	0.29	0.3	0.76	0.47
RMSE	0.49	0.52	0.48	0.55
Accuracy	75%	73%	77%	70%

Table 1: Inter-rater agreement and accuracy of our prediction models.

Results

We asked 646 workers to rate 2,000 Twitter accounts, all of which tweeted in English. Each Twitter account was annotated, on average, 3.12 times. For each account rated, we asked the worker to provide their opinion regarding the following traits of the profile owner: gender, age range (18-, 18-25, 25-30, 30-45, 45+), life stage (high school, university, young professional, established professional, retired). The worker was additionally asked to rate the following traits of the profile owner on a five points scale (very low to very high): tech-savviness, knowledge level of computer games, trustworthiness, quality and depth of generated content, humor, originality, level of influence. We also asked whether the profile owner is likely to be a fan of the Xbox platform or of the Playstation platform.

We use two measures for the quality of prediction. One is the root mean square error (RMSE) of the numerical predictions (on the five point scale or on a $\{0, 1\}$ Boolean scale). The second is based on partitioning the user population into thirds, by sorting the users from the highest to lowest score of the predicted trait. We can then examine the two extreme thirds, and check the prediction accuracy of determining whether a user is in the top or bottom third (ignoring the middle third). Table 1 shows the accuracy of our predictions (measured using 10-fold cross validation).

Table 1 indicates that it is indeed possible to predict many perceptions about gamers from the language they use in online social networks. The table also lists the agreement score for each variable (ICC). As the table indicates, some properties are more difficult to determine than others.

We have also built similar models to predict the time a gamer has spent playing computer games and their actual performance in computer games.[2] Our results indicate a prediction accuracy 59% for the time spent playing games, and 64% for actual performance in playing games, referred to as "Gamer Score" (the accuracy is for separating users in the top third and those in the bottom third of these properties).[3] This indicates that information from online social networks is at least somewhat predictive of various objective gaming related traits of users.

[2] We had access to such data from the user profiles in the Xbox platform data.

[3] Interestingly, our methods achieved better predictions for the actual ability in playing computer games than for the time spent playing them.

References

Bachrach, Y., and Porat, E. 2015. Fingerprints for highly similar streams. *Information and Computation*.

Bachrach, Y.; Kosinski, M.; Graepel, T.; Kohli, P.; and Stillwell, D. 2012. Personality and patterns of Facebook usage. In *Proceedings of WebSci*, 24–32.

Bachrach, Y.; Ceppi, S.; Kash, I. A.; Key, P.; Radlinski, F.; Porat, E.; Armstrong, M.; and Sharma, V. 2014a. Building a personalized tourist attraction recommender system using crowdsourcing. In *AAMAS*, 1631–1632.

Bachrach, Y.; Finkelstein, Y.; Gilad-Bachrach, R.; Katzir, L.; Koenigstein, N.; Nice, N.; and Paquet, U. 2014b. Speeding up the xbox recommender system using a euclidean transformation for inner-product spaces. In *RecSys*.

Bachrach, Y.; Graepel, T.; Kohli, P.; Kosinski, M.; and Stillwell, D. 2014c. Your digital image: factors behind demographic and psychometric predictions from social network profiles. In *AAMAS*, 1649–1650.

Bachrach, Y. 2015. Human judgments in hiring decisions based on online social network profiles. In *DSAA*.

Clements, M.; Serdyukov, P.; de Vries, A. P.; and Reinders, M. J. T. 2011. Personalised travel recommendation based on location co-occurrence. *CoRR* abs/1106.5213.

Golbeck, J.; Robles, C.; Edmondson, M.; and Turner, K. 2011. Predicting personality from Twitter. In *Proceedings of SocialCom/PASSAT*.

Koren, Y.; Bell, R.; and Volinsky, C. 2009. Matrix factorization techniques for recommender systems. *Computer* (8):30–37.

Kosinski, M.; Stillwell, D.; Kohli, P.; Bachrach, Y.; and Graepel, T. 2012. Personality and website choice.

Kosinski, M.; Bachrach, Y.; Kohli, P.; Stillwell, D.; and Graepel, T. 2014. Manifestations of user personality in website choice and behaviour on online social networks. *Machine Learning* 95(3):357–380.

Kosinski, M.; Stillwell, D.; and Graepel, T. 2013. Private traits and attributes are predictable from digital records of human behavior. *National Academy of Sciences*.

Lewenberg, Y.; Bachrach, Y.; and Volkova, S. 2015. Using emotions to predict user interest areas in online social networks. In *DSAA*.

Schwartz, H. A.; Eichstaedt, J. C.; Kern, M. L.; Dziurzynski, L.; Ramones, S. M.; Agrawal, M.; Shah, A.; Kosinski, M.; Stillwell, D.; Seligman, M. E.; et al. 2013. Personality, gender, and age in the language of social media: The open-vocabulary approach. *PloS one* 8(9):e73791.

Ustinovsky, Y., and Serdyukov, P. 2013. Personalization of web-search using short-term browsing context. In *Proceedings of the 19th ACM International Conference on Information and Knowledge Management (CIKM)*, 1979–1988.

Yang, W.-S.; Dia, J.-B.; Cheng, H.-C.; and Lin, H.-T. 2006. Mining social networks for targeted advertising. In *System Sciences, 2006. HICSS'06. Proceedings of the 39th Annual Hawaii International Conference on*, volume 6, 137a–137a. IEEE.