

# Better the Tweeter you know: social signals on Twitter

M.J. Chorley\*, G.B. Colombo<sup>†</sup>, S.M. Allen<sup>‡</sup>, R.M. Whitaker<sup>§</sup>  
School of Computer Science & Informatics  
Cardiff University, CF24 3AA, UK  
{m.j.chorley\*, g.colombo<sup>†</sup>, stuart.m.allen<sup>‡</sup>, r.m.whitaker<sup>§</sup>}@cs.cf.ac.uk

**Abstract**—We present results from a web-based experiment conducted to assess the effect of Twitter metadata on decision making in content consumption. Participants were presented with information concerning two tweets and asked which they would prefer to read. Analysis of the results shows that recognition of the author as being within the readers local network is highly influential in the decision to read a tweet. This has analogies with results from cognitive psychology on decision making processes such as the recognition heuristic. The role of more detailed quantitative metadata has also been assessed. Surprisingly, metadata describing the popularity of tweet authors in terms of the number of followers or the number of tweets written has no significant impact on decision making, while metadata describing the tweet content (the number of retweets) has a significant impact, with a large proportion of users preferring to read content that has been retweeted a larger number of times. When friendship information and quantitative values are combined the impact of the friendship information is reduced, but a larger proportion of users still prefer to choose based on this information, while the impact of the retweet value is reduced.

## I. INTRODUCTION

The real-time nature of micro-blog services such as Twitter<sup>1</sup> leads to a constantly updating stream of content whose entire consumption can require a significant cognitive effort. Thus when reading/browsing Twitter humans perform a subconscious filtering process through which decisions for consumption are made. Although quick glimpses of parts of the actual tweet text can contribute to users choices (through noticing items such as hashtags or notable keywords), other metadata cues external to the content of a tweet also influence the selection mechanisms of readers. For example, tweets may be perceived as being more worthy of attention when their author is recognised as being within a reader's social circle, irrespective of content. Metadata cues are also interesting because they are key elements in exposing readers to unexplored, yet relevant social media content. It is not sufficient to merely display such content; readers must also be provided with appropriate cues that avoid them skipping, dismissing or ignoring the content.

In this paper we investigate the role of such metadata as cues for assessing relevance, and as such, our work is closely related to cognitive decision making under constrained conditions. In particular, we are interested in determining the

extent to which simple psychological models such as the Recognition heuristic [1] apply within the context of tweets and Twitter users. The Recognition heuristic states: “*If one of two objects is recognised and the other is not, then infer that the recognised object has the higher value with respect to the criterion.*” These cognitive approaches for decision making assume that cues which are based on familiarity drive human preference. For example, in the original experiments [1] a number of participants were asked to choose which from a group of German cities had the highest population with the results showing that they routinely (and correctly) picked the city they recognised. To investigate these issues we have developed an open online experiment based on the pairwise comparison of selected tweets. A Twitter user is asked to make choices on their preference of tweet for consumption when they are presented with only limited meta-data. In each pair of tweets presented to a user, one is selected from their timeline (the list of tweets they would personally see when browsing twitter.com, written by the people they follow) and one comes from a user whom they definitely do not follow, thus being a tweet they would not normally see. We present users with limited information about each tweet, but do not show the content itself, and ask them which from the pair they would prefer to read. We enforce that the participants decision is taken on explicit cues, either qualitative information (such as the authors screen name) or quantitative information (such as the number of retweets of the tweet).

The remainder of the paper is structured as follows: Section II gives an overview of the related literature; Section III provides details of the experimental design, while Section IV presents and analyses the results obtained. Finally Section V summarises the conclusions of the work and outlines future enhancements.

## II. RELATED WORK

Micro-blogging services have seen a remarkable growth in the last few years, partly due to the limited cognitive effort required to parse an individual update in return for the numerous benefits that they can provide. These services are used for multiple purposes from social networking to advertising; from receiving and broadcasting news feeds to exchanging information targeting specific topics or communities. One of the reasons for their success is the opportunity to post and receive updates in real time so to draw attention events while

<sup>1</sup><http://www.twitter.com>

they are occurring [2].

Exploiting the fact that Twitter updates often contain links to URLs, [3] shows that this can be used to provide a ranking of web pages. In particular, it is emphasised that micro-blogs can allow searches to integrate currently popular trends by adding the most recent links to the search results; this feature is referred to as ‘recency ranking’.

In [4] an algorithm based on PageRank is proposed, in which web links are simply replaced by the follower and following relationships in Twitter. In addition, machine learning techniques based on ‘bag of words’ assumptions are used to further filter updates by topical similarity [4], thus exploiting the underlying property of ‘homophily’ in social networking (i.e. users tend to make social links to others with similar preferences, behaviour, and interpersonal characteristics). [5] points out that mechanisms based on link counting all tend to weight links in a uniform way. The authors propose a Twitter user ranking in which tweets’ scores are represented by the number of re-tweets they have and users increase their authority when posting highly retweeted updates and/or by being followed by other authoritative users.

The ranking of users by quality and topical relevance has largely focused on the ‘following/followers’ relationship only [6]. This work then proposes a variant of PageRank based on the implicit retweet sub-graph (i.e. links between users are placed if they retweet each others update at least once, with weight defined accordingly). The authors show that the transitivity of topical relevance is better preserved over retweet links, and that retweeting a user is a significantly stronger indicator of topical interest than following him.

As with other works described in this section, it appears that the cognitive cues that make tweets by specific authors more attractive to others (e.g. ranked highly) lie within the graphs defining follower and following relationships as well as the quality of the posts written (that can be quantified as proportional to the number of re-tweets of a given update) [7]. In addition, there is the social components of including friends (also from other social networks) in personalised ‘reading lists’ (also organised by interest preferences) [6]. However, deeper investigations aiming to compare and weight different cues have yet to be produced.

A very recent article [8] gives a qualitative description of some of these cues focusing on those defined as ‘credibility cues’ (i.e. to evaluate if a certain tweet is true or false). These are outlined as: ‘author influence’ (measured by follower, retweet, and mention counts); ‘topical expertise’ (from information given by the Twitter user profile); ‘history of on-topic tweeting’ (also considering pages outside of Twitter and updates with location relevant to the particular topic); ‘user’s reputation’ (the author is someone a user follows or has heard of, or is given additional verifications from the Twitter website). In contrast to this we investigate the influence of metadata cues and the role that they play as a filtering mechanism for establishing priority of interest. This is useful for further understanding human behaviour but also provides a basis for system design, search and automated filtering.

### III. EXPERIMENTAL DESIGN

We have designed an experiment called TweetCues to provide evidence on the cues used by humans when deciding the importance of tweets. This is an example of constrained decision making, here models such as the Recognition heuristic are relevant [1]. Applying this heuristic to Twitter suggests that if a user is presented with information about two tweets but not their content, and asked to choose which they wish to read, then they will choose the tweet they recognise as coming from their own timeline over a tweet from an author that they do not recognise. The experiment will further investigate this effect by considering the relative importance of recognition compared to quantitative measures that may be used to explicitly rank the importance of an unseen message. In particular, we aim to derive insight into the following questions:

- Q1 *In the absence of any further information, do participants prefer tweets that may be recognised as coming from their personal timeline?*
- Q2 *Which quantitative metrics do participants use as cues to determine the quality of an unseen tweet?*
- Q3 *Do the quantitative metrics or recognition have the larger effect on the preferences of the participants?*

The experiment has been implemented as a web application which allows participants to log in using their Twitter credentials, after which they are asked a series of questions to assess their use of certain cues when judging the value of unseen information. A subject is invited to take part in the experiment by accessing the website and is required to log in to the website using their Twitter credentials in order to access information from their Twitter account. To ensure genuine Twitter users with sufficient tweets in their timeline, subjects are only permitted to continue to the experiment if they have at least 10 followers and are following 10 others. The participant is subsequently presented with a series of questions which display partial information about a pair of tweets. In each pair, one of the tweets is chosen from the user’s timeline, while the other is chosen at random from a pool of tweets stored by the application which are known to have no connection to the user. A subset of the associated metadata is displayed for each of the tweets, and they are asked the following question: ‘Which tweet would you like to read?’. The participant then selects the tweet that they would like to read, it is displayed to them, and the process is repeated in succession with different combinations of metadata being exposed. To reduce the effect of ‘automatic’ selection of answers (for example, automatically clicking on all of the left items or answering without paying attention to the questions) the order of the tweets displayed is chosen randomly for each question. This selection bias is also limited by the user being required to click a centrally displayed ‘continue’ button between each question, so resetting the cursor position on screen for each question.

One of the main drawbacks of internet based experiments is a higher dropout rate than in normal lab experiments [9]. However, this obviously depends on experiment specific issues

such as its duration and the cognitive effort required. The duration of our experiment is limited to a few minutes (we recorded an average completion time of around 4 minutes). In our experiment, we observed a completion rate of 78.9% from 113 participants.

### A. Questions and InfoTypes

In order to separate the influence of different cues on the user, we define a number of “InfoTypes” to specify which information about a tweet is shown to a participant in any one question. These can be partitioned into two distinct subgroups:

- *Friendship*

The first subgroup contains information related to the identification of the author, either directly (such as providing their friendship relationship with the author), or indirectly (for example by displaying their profile image or their screen name). Information in this group includes: *Screen name*, *Name*, *Avatar*, and *Friendship*. Note that information about friendship is either directly or indirectly revealed (with different degrees of recognition).

- *Quantitative*

The second subgroup contains quantitative measures characterising either the author within the Twitter network, or the tweet itself. These are *Follower count*; *Following count* (giving the number of followers of the author or the number of others that he/she follows); *Tweets count* (the number of updates posted by the author); and *Number of Retweets* (for the given update).

We define ‘Question Types’ to allow us to control InfoTypes presented in an individual instance of the survey. To allow a correct statistical interpretation of the data we have selected 25 key InfoType combinations as the most relevant for the scope of this research. Each user receives the same number of questions and each of the considered combinations is shown precisely once within the same survey. These combinations are listed in Table I. Note that the questions involve ‘single cues’ (to explore the impact of an individual InfoType, see Figure 1) and ‘combined cues’ questions (to investigate the impact of aggregated cues, see Figure 2). In order to minimise the effect of drop-outs from the survey (to ensure we have at least some responses for each InfoType even if participants do not complete all questions) we display a permutation of the twenty five possible InfoType combinations according to a uniform random distribution; that is the order in which the questions are presented is randomly chosen. The actual number  $N$  of answers received for each question is shown in Table I.

## IV. EXPERIMENTAL RESULTS

When considering proportions of users selecting a particular tweet we examine the variables ‘selection of an item from a user’s timeline’ and ‘selection of the item with the greatest numerical value’, herein described as *Timeline* and *Greatest*. When considering quantitative information, questions where equal values are presented to the user are discarded during analysis.



Fig. 1. Single Cue question - only one InfoType shown to participant - pre selection state

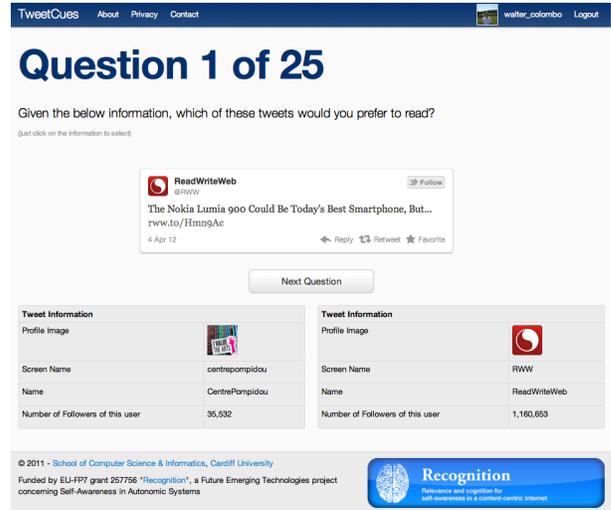


Fig. 2. Combined Cues question, multiple InfoTypes shown to participant - after selection state

TABLE I  
LIST OF QUESTION TYPES

#	N	Question Type
1	107	Screen Name
2	113	Name
3	109	Avatar
4	109	Friendship
5	107	Screen Name, Name and Avatar
6	106	Follower Count
7	109	Following Count
8	107	Tweet Count
9	106	Number of Retweets
10	109	Screen Name and Follower Count
11	109	Screen Name and Following Count
12	109	Screen Name and Tweet Count
13	110	Screen Name and Number of Retweets
14	110	Avatar and Follower Count
15	106	Avatar and Following Count
16	108	Avatar and Tweet Count
17	108	Avatar and Number of Retweets
18	108	Friendship and Follower Count
19	108	Friendship and Following Count
20	107	Friendship and Tweet Count
21	113	Friendship and Number of Retweets
22	111	Screen Name, Name, and Avatar and Follower Count
23	107	Screen Name, Name, and Avatar and Following Count
24	113	Screen Name, Name, and Avatar and Tweet Count
25	112	Screen Name, Name, and Avatar and Number of Retweets

TABLE II  
PROPORTION OF USERS SELECTING PARTICULAR TWEETS FOR EACH OF  
THE ‘SINGLE CUE’ QUESTION TYPES

QuestionType	InfoType	proportion	
		$P_T$	$P_G$
1 - Screen Name	Friendship	<b>0.89719</b>	-
2 - Name		<b>0.88495</b>	-
3 - Avatar		<b>0.88073</b>	-
4 - Friendship		<b>0.88073</b>	-
5 - Names+Avatar		<b>0.85981</b>	-
6 - Followers	Quantitative	0.43396	0.54716
7 - Tweets		0.50458	0.50458
8 - Following		0.53271	0.50467
9 - Retweets		<b>0.37735</b>	<b>0.90425</b>

### A. Single Cue Questions

Table II reports the proportion of subjects selecting the tweet from within their timeline ( $P_T$ ) and the proportion selecting the greatest value ( $P_G$ ) respectively. There is a striking difference between the friendship and quantitative questions. As may be expected, in over 85% of cases the participant selected the ‘recognised’ tweet from their timeline when given identifying information about the author. For each individual item of identifying information, the participants chose the tweet from their timeline significantly more often ( $\chi^2$  tests with  $p < 0.0001$ ). For the quantitative questions, the results show that there is no significant effect based on the number of followers, friends or tweets ( $\chi^2$  tests with  $p > 0.3314$ ), but that information on retweets is significant ( $\chi^2$  test with  $p < 0.0001$ ). Other than with retweets, it appears that the selection of tweets is random, consistent with the results showing tweet with the highest value is chosen in between 50% to 55% of cases. As can be seen from the proportion of these taken from the timeline (between 43% and 54%), it is unlikely that participants are basing their decisions on any implicit information they have about their local network. This indicates that when the author of the tweet is not known, their popularity and prolificacy have little influence. Instead, it is only the one metric that applies directly to the popularity of the content itself (the number of times it has been forwarded) that has an impact on the participants decisions.

We examine whether the participant made their choice with a consistent strategy relative to the values displayed (always choosing the lowest value or always choosing the highest value) across all three single cue questions in which no preference for the greatest or lowest value is seen: those showing the number of followers, the number the author is following, and the number of tweets made. Analysis shows that in these questions users are not selecting with any consistent strategy. Only 23% of users acted consistently (choosing either all the greatest or all the smallest values). This is consistent with a random selection, as with 8 possible combinations of choice across the three questions, only two would result in a consistent selection.

Note that, although at first the results concerning quantitative values look surprising, it is in line with a number of published works [10], [5], [6], [4] concerning tweet ranking

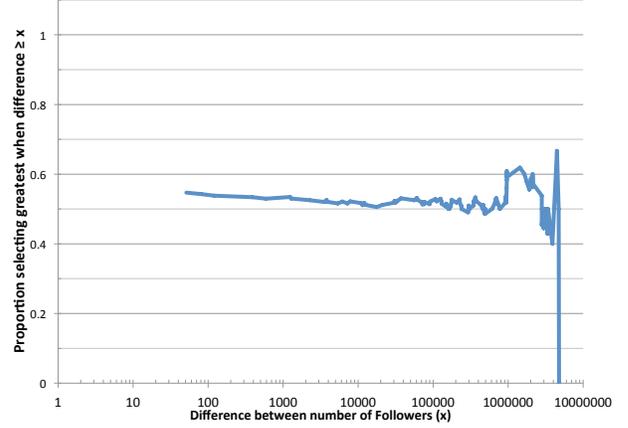


Fig. 3. Proportion selecting the greatest for difference between number of followers

that also criticise how some of the most used ranking methods tend to give priority to personal characteristics of the authors in relation to their position in the Twitter network (i.e. they are given more ‘authority’ when having a large number of followers, people followed, or total posts written), whereas real information about the quality of tweets is not really considered. In addition, these works show how including the ‘number of retweets’ as a quality measure can considerably improve the performance of the methods proposed.

### B. Dominance of High Quantitative Values in Single Cue Questions

We have seen that, except for retweets, for all single cue quantitative questions there is no evidence that the information displayed is used by the participant when making a decision. However, it might be expected that such decisions are affected by the relationship between the two quantitative values presented. It may be expected for instance that when the difference between the two values is small (both tweet’s authors have similar follower counts for instance), participants behave randomly, choosing either of the smallest or largest value, but that when the difference is large (so one tweet author is much more popular than the other), the behaviour is more ordered. However, examining the probability distribution in Figure 3 suggests that this is not the case for the followers metric. Here, the proportion of questions in which the tweet with the greatest value was selected is shown for all differences higher than some value  $x$ . It is clear to see that this proportion does not change significantly as the magnitude of this difference increases (the noise at the higher range can be attributed to the decrease in sample size at these values). The proportion holds steady at around 0.5 for all differences up to around 1,000,000 followers. There is no sign of users using a different selection mechanism when the difference between the number of followers presented is very large or very small.

If we examine the same proportions for the number of retweets (Figure 4), where we saw a significant impact on the

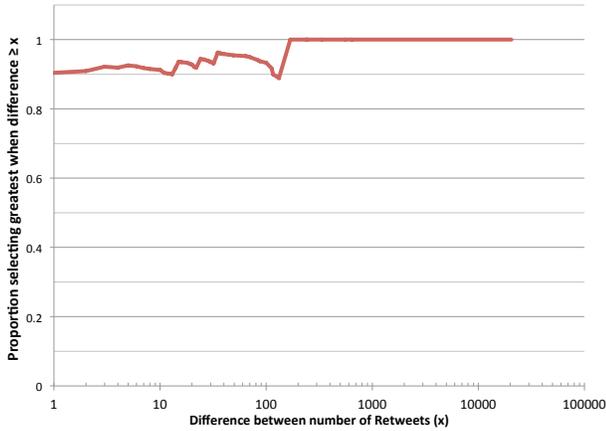


Fig. 4. Proportion selecting the greatest for difference between number of retweets

TABLE III  
PROPORTION OF USERS SELECTING THE TIMELINE TWEET ( $P_T$ ) FOR EACH OF THE ‘COMBINED CUES’ QUESTION TYPES

Screen name	Followers	Following	Tweets	Retweets
Avatar	0.779	0.816	0.798	0.745
Friendship	0.772	0.820	0.861	0.731
Friendship	0.814	0.738	0.775	0.654
Names +Avatar	0.846	0.822	0.929	0.848

decision making of participants, we see that the proportion selecting the greatest value remains high across all difference values. This shows that even when the difference between the number of retweets of two tweets is very small, there is still a high probability that a participant will choose the greatest value.

### C. Combined Cue Questions

Table III and IV presents the proportion of users selecting respectively the timeline ( $P_T$ ) and the greatest tweets ( $P_G$ ) for the ‘combined cues’ questions, each of which combines a friendship cue with a quantitative measure. We observe that the mean values concerning the selection of a tweet from the timeline remains close to 0.8, whereas no particular tendency is directly observed for the selection of the greatest (values range from 0.38 to 0.67). These results suggest that participants are more likely to use the friendship cues to select an update that comes from their own timeline, ignoring the quantitative cues (no matter what their value). From this table we can see that in all cases the impact on subjects of the friendship cues to direct their selections is weaker than in the single cue experiments (since the proportions of users selecting the tweet from within their timeline are generally lower).

To further investigate this we conducted a one-way ANOVA on the statistical significance of the differences between combined cues questions and the corresponding single cue questions, shown in Table V. Here sources of variation are represented by pairs of single versus combined cues. The random variables on which we conduct the analysis are the

TABLE IV  
PROPORTION OF USERS SELECTING THE GREATEST TWEET ( $P_G$ ) FOR EACH OF THE ‘COMBINED CUES’ QUESTION TYPES

	Followers	Following	Tweets	Retweets
Screen name	0.449	0.472	0.522	0.615
Avatar	0.427	0.514	0.472	0.536
Friendship	0.416	0.467	0.401	0.674
Names +Avatar	0.441	0.383	0.495	0.500

TABLE V  
ONE WAY ANOVAS FOR DIFFERENT ‘COMBINED CUES’ QUESTIONS

Source of var-QT		X var.	Mean Square		F statistic	
Friendship	Quantity		betw.	with.	calc.	tab.
Screen name	Followers	Timeline	0.672	0.059	<b>11.24</b>	< 3.92
	Following	Timeline	0.378	0.057	<b>6.63</b>	< 3.92
	Tweets	Timeline	0.565	0.057	<b>9.89</b>	< 3.92
	Retweets	Timeline	1.051	0.062	<b>16.84</b>	< 3.92
Avatar	Followers	Timeline	0.775	0.059	<b>13.13</b>	< 3.92
	Following	Timeline	0.293	0.053	<b>5.49</b>	< 3.92
	Tweets	Timeline	0.073	0.050	1.46	> 3.84
Friendship	Retweets	Timeline	1.486	0.063	<b>23.25</b>	< 3.92
	Followers	Timeline	0.454	0.061	<b>7.46</b>	< 3.92
	Following	Timeline	1.163	0.067	<b>17.29</b>	< 3.92
Names+Avatar	Tweets	Timeline	0.654	0.063	<b>10.29</b>	< 3.92
	Retweets	Timeline	2.618	0.074	<b>34.99</b>	< 3.92
	Followers	Timeline	0.727	0.049	1.45	> 3.84
Retweets	Following	Timeline	1.592	0.051	3.07	> 3.84
	Tweets	Timeline	0.070	0.036	1.95	> 3.84
	Retweets	Timeline	0.110	0.049	2.21	> 3.84
Retweets	S.name	Greatest	3.836	0.086	<b>44.10</b>	< 3.92
	Avatar	Greatest	6.084	0.089	<b>67.96</b>	< 3.92
	Friend.	Greatest	2.547	0.083	<b>30.49</b>	< 3.92
	Nam.+Av.	Greatest	7.600	0.089	<b>84.69</b>	< 3.92

‘selection of the tweets inside the subject timeline’ and the ‘selection of the tweets with greatest value’; the ‘within group’ values refer to the total number of observations (i.e answers received for a given Question Type). For space restrictions we will only focus on two specific groups of questions.

The first group concerns cues about friendship and includes all ‘combined’ question types from number 10 to 30 questioning whether the addition of a quantitative cue weakens (as may be expected) the impact on the ‘selection of the timeline update’ (considered as the random variable). Rejecting the null hypothesis of no such impact (in bold in the table) means that users decision to choose the timeline tweet are significantly weakened because of the addition of the quantitative cue (and vice versa for the values not highlighted in bold). As can be observed from the results this happens in most cases, showing that the difference is significant. Note that when multiple friendship information is provided (i.e. name, screen name, and avatar, for questions 22 to 25) and in one case when the avatar is shown in combination with the tweet count (question 16), the addition of the quantitative cue does not have a significant negative effect on the impact of the friendship cues. This result is not completely unexpected since showing images or multiple information could be thought as having stronger impact on participants than showing only the author names.

If we look at quantitative cues, the only cue that had significant impact in isolation (see the single cue Table II) was the ‘number of retweets’. We can consider with interest

the questions that combine this information with friendship information (questions 13, 17, 21, 25). In all these cases the null hypothesis of no impact of the additional friendship cues is always rejected. In other words, the addition of friendship information to these questions has significantly affected the decision making process of the participants. Also from Table III in these cases involving the display of the number of retweets the value of  $P_G$  drops towards 0.5 that shows no significant impact of the ‘retweet’ cues in the selection. Thus we can conclude that the addition of friendship information changes the process of decision making in the cases where the number of retweets is shown. However, this process is still skewed slightly by the number of retweets being displayed, as evidenced by the slightly higher values for  $P_G$  in the relevant questions in Table IV.

## V. CONCLUSION AND FUTURE WORK

An online user-based experiment has been conducted that aims to analyse the effects of various tweet meta-data on the decision making process of users as it applies to deciding which content to consume. Analysis of the experimental results has answered several questions (as originally stated in Section III).

Firstly (Q1), we have shown that in the absence of any further information participants prefer tweets that can be recognised as coming from their personal timeline. A significant proportion of users will select the tweet that is recognised over one that is not. This demonstrates that the Recognition heuristic for decision making operates as may be expected with Twitter metadata. When asked to make constrained decisions, participants prefer the object they recognise to the one that they do not.

Secondly (Q2), we have seen that when viewed individually only one quantitative value has an effect on the decision making process: the number of retweets. Information related to popularity or status of the authors of tweets is not judged as important, with users selecting randomly when presented with either the number of followers, the number of following or the number of tweets written by the author. The metric relating to the content itself is however viewed as important, with a significant proportion of users selecting the greatest retweet value as the content they would like to consume. This is in contrast with some of the current tweet ranking algorithms that prioritise metadata describing the authority of authors rather than the content of the updates.

Thirdly (Q3), we have seen that the qualitative friendship metadata related to recognition has a stronger effect on the preferences of participants than the quantitative information. Results suggest that when presented with multiple types of information, a larger proportion of users tend to select tweets based on the recognition of an author from their timeline than select based on a larger quantitative value. However, this pro-

portion of users selecting based on recognition is significantly reduced in most cases by the addition of quantitative values. Interestingly, the only quantitative value that has an effect when viewed individually (the number of retweets) is also affected significantly when additional friendship information is supplied. In this case, the recognition effect becomes more important, and the number of retweets is no longer a significant factor in decision making.

Future plans include the application of the presented experiment for implicit provision of Twitter updates. The effect of the studied metadata on the decision making process can be used in future content provision systems to improve user satisfaction with provided content. As previously mentioned binary comparison trials are recognised as the most successful model for ranking items within very large groups (under the condition that items to compare are selected giving higher weight to the most highly rated items [10]).

## ACKNOWLEDGMENT

This research has been funded by RECOGNITION grant 257756, an EC - FP7 Future Emerging Technologies project concerning Self-Awareness in Autonomic Systems. The authors gratefully acknowledge the advice of Marc Buehner of Cardiff University School of Psychology for his assistance in experiment design and analysis.

## REFERENCES

- [1] D. G. Goldstein and G. Gigerenzer, “Models of ecological rationality: the recognition heuristic.” *Psychological review*, vol. 109, no. 1, pp. 75–90, Jan. 2002.
- [2] A. Java, X. Song, T. Finin, and B. Tseng, in *Proc. of the 9th WebKDD and 1st SNA-KDD 2007 workshop on Web mining and social network analysis*. New York, NY, USA: ACM, pp. 56–65.
- [3] A. Dong, R. Zhang, P. Kolari, J. Bai, F. Diaz, Y. Chang, Z. Zheng, and H. Zha, “Time is of the essence: improving recency ranking using twitter data,” in *Proc. of the 19th international conference on World wide web*, ser. WWW ’10. New York, NY, USA: ACM, 2010, pp. 331–340.
- [4] J. Weng, E. P. Lim, J. Jiang, and Q. He, “TwitterRank: finding topic-sensitive influential twitterers,” in *Proc. of the third ACM international conference on Web search and data mining*, ser. WSDM ’10. New York, NY, USA: ACM, 2010, pp. 261–270.
- [5] Y. Yamaguchi, T. Takahashi, T. Amagasa, and H. Kitagawa, “Turank: twitter user ranking based on user-tweet graph analysis,” in *Proc. of the 11th international conference on Web information systems engineering*, ser. WISE’10. Berlin, Heidelberg: Springer-Verlag, 2010, pp. 240–253.
- [6] M. J. Welch, U. Schonfeld, D. He, and J. Cho, “Topical semantics of twitter links,” in *Proc. of the fourth ACM international conference on Web search and data mining*, ser. WSDM ’11. New York, NY, USA: ACM, 2011, pp. 327–336.
- [7] E. Rocha, A. Francisco, and H. Calado, P. and Sofia-Pinto, “User profiling on Twitter,” *Semantic Web*, submitted 2011.
- [8] A. Phelps, “Think fast: Is that tweet true or false? how we use credibility cues to make decisions,” *Nieman Journalism Lab*, 2012. [Online]. Available: <http://www.niemanlab.org/>
- [9] U. Reips, “Internet-based psychological experimenting: five dos and five don’ts,” *Soc. Sci. Comput. Rev.*, vol. 20, no. 3, pp. 241–249, Sep. 2002.
- [10] A. Das Sarma, A. Das Sarma, S. Gollapudi, and R. Panigrahy, “Ranking mechanisms in twitter-like forums,” in *Proc. of the third ACM international conference on Web search and data mining*, ser. WSDM ’10. New York, NY, USA: ACM, 2010, pp. 21–30.