

# Taking It All In? Visual Attention in Microblog Consumption

Scott Counts<sup>1</sup> & Kristie Fisher<sup>2</sup>

<sup>1</sup>Microsoft Research, One Microsoft Way, Redmond, WA 98051

<sup>2</sup>University of Washington, Department of Psychology, Box 351525, Seattle, WA 98195

counts@microsoft.com, kjfisher@u.washington.edu

## Abstract

Microblogging environments such as Twitter present a modality for interacting with information characterized by exposure to information “streams”. In this work, we examine what information in that stream is attended to, and how that attention corresponds to other aspects of microblog consumption and participation. To do this, we measured eye gaze, memory for content, interest ratings, and intended behavior of active Twitter users as they read their tweet streams. Our analyses focus on three sets of alignments: first, whether attention corresponds to other measures of user cognition such as memory (e.g., do people even remember what they attend to?); second, whether attention corresponds to behavior (e.g., are users likely to retweet content that is given the most attention); and third, whether attention corresponds to other attributes of the content and its presentation (e.g., do links attract attention?). We show a positive but imperfect alignment between user attention and other measures of user cognition like memory and interest, and between attention and behaviors like retweeting. To the third alignment, we show that the relationship between attention and attributes of tweets, such as whether it contains a link or is from a friend versus an organization, are complicated and in some cases counterintuitive. We discuss findings in relation to large scale phenomena like information diffusion and also suggest design directions to help maximize user attention in microblog environments.

## Introduction

As people increasingly use social media channels as sources of information, explicit information seeking behaviors such as search are getting augmented by a new form of information consumption characterized by exposure to a “stream” of information. These streams, such as a person’s tweet stream in Twitter, often are defined globally by user-specified parameters such as the set of people a person follows, but the actual content is supplied by other people and is unknown to the user until she encounters it. Thus in the most proximal sense, what pieces of information in this new modality impact the user are

determined by the user’s attention: what in the information stream does the user actually attend to?

This critical attention process characterizes the very moment when “the information meets the user.” That is, without first attending to a piece of information, the user cannot take action on the information or perform behaviors that play a role in a number of important network-level processes. Information diffusion, for example, cannot happen without users first attending to pieces of information in order to decide whether or not to pass them along in the network. Likewise, linked content cannot be explored without the user first attending to the surrounding tweet or status update to decide whether the content is worth clicking through to.

In this paper we explore this question of what users attend to in social media streams with respect to three sets of alignments. The first concerns user cognition: does what the user attends to align with what she finds interesting and what she remembers later? The second concerns user behavior: does what the user attends to align with behaviors like retweeting and replying? The third concerns the social media environment: does what the user attends to align with the attributes and presentation of the information? These alignments are important because of the role they play in the efficiency with which network level processes occur. Taking again information diffusion as an example: to the extent users are attending to information in social media streams in a suboptimal way, diffusion will be inefficient or will mischaracterize true user interest. That is, the user may retweet item B when in fact she would have chosen instead to retweet item A had her attention been properly drawn to it.

Our research goal, then, is to quantify these alignments. Taking Twitter as our information source, we use eye tracking to measure visual attention while reading tweets. We then examine how this visual attention corresponds to interest and memory as well as to properties of the tweets, such as whether they are from friends versus organizations, are retweets, or contain links.

## Background

### The Microblogging Information Stream

Microblogging environments like Twitter are used by millions of people to share and consume short bits of information. Although there is likely a wide range of variability among users, a recent estimate suggests that an active user may see over 700 tweets in her daily stream (Bernstein et al., 2010). Research on Twitter has shown that it can be a news source for users (Kwak et al., 2010), that aspects of Twitter's social structure are conducive to information diffusion (Yang and Counts, 2010b), and that tweet posting rates can be extremely fast (Yang and Counts, 2010a). In short, Twitter can serve as a very large and rapidly-changing information source.

As users consume these dynamic information streams, they occasionally interact with them in order to click a shared link, reply to someone, or retweet another user's tweet. Because these behaviors are measurable and largely publicly available, much research on Twitter has used them as a primary assessment tool (e.g., Suh et al., 2010; Yang and Counts, 2010b). Yet we also know that the set of items for which users take no action is likely the overwhelming majority of items. For instance, boyd et al. (2010) report that only 3% of tweets are retweeted. Thus, research focusing on this small subset of tweets is likely not accurately capturing a user's full experience.

### Human Information Processing and Twitter

Core principles of how people process new information have been established by decades of previous work on memory, learning, and problem solving. First, a person must attend to a piece of new information so that it can be encoded into working memory (see Baddeley 1997 for a detailed overview of memory theory). The most effective encoding involves a "deep" processing wherein meaningful connections are made among the pieces of new information and to one's existing knowledge base ( Craik and Lockhart, 1972). Such processing requires time and mental effort.

The nature of Twitter presents several challenges to performing such "deep" information encoding effectively. Tweets are short text strings and any other media must be linked to text in the tweet. Furthermore, all of these small pieces of information from multiple senders are presented with little meaningful organization. Deep encoding may also be hampered by time constraints, and some tweets may not contain interesting or useful content. Therefore, if users want to consume useful information from Twitter as efficiently as possible, they must first evaluate the interest level of each tweet and then effectively encode the content from the relevant tweets, all within a limited timeframe. Given this cognitively challenging task, it is not surprising that allocation of attention within an information stream like Twitter might be inefficient.

## Eye Tracking

To investigate attention in the context of microblogging environments such as Twitter, we turned to eye-tracking. Eye gaze metrics have been assumed to represent the most outward demonstration of cognitive processes, and as such they can provide an index of attention capture and information encoding (Just and Carpenter, 1976). In the domain of HCI, previous work on web browsing (e.g., Buscher et al., 2009), web search (e.g., Cutrell and Guan, 2007), and information foraging (e.g., Chi, Gumbrecht, & Hong, 2007) has demonstrated that eye-tracking can be used to assess where users direct their attention on a web page. In our study, using eye-tracking we sought to assess what attributes of tweets capture users' attention, how long tweets can hold attention, and how effectively users direct their gaze when reading through their Twitter feed.

## Research Questions

Our specific research questions were as follows:

- (RQ1) Visual attention for tweets with respect to other measures of use cognition:
  - How long do users spend reading each tweet?
  - How much is actually remembered by users?
  - How much do users consider highly interesting?
- (RQ2) How well do behaviors (replying, retweeting), correspond to attention?
- (RQ3) Do particular properties of tweets (e.g., hashtags, links) better capture attention or otherwise affect how users consume their content?

## Methods

To investigate these research questions, we recruited active Twitter users and tracked their eye movements as they read their Twitter feed. Participants then completed a series of questionnaires, including a memory task, about the Tweets they saw, as described in detail below.

### Participants

Twenty people (4 females) from our company participated in our study (MAge = 30.21 years, SDAge = 8.93); they were recruited via e-mails sent to several listservs. Participants were required to have at least 50 friends on Twitter, to regularly check their Twitter feed at least 4 times a week, and to have posted at least 10 tweets since joining Twitter. Note that while the sample size is relatively small, it is not uncommon for eye tracking studies (e.g., Buscher et al, 2009). Also, while the gender balance was more male-skewed than we would have liked, we are not aware of any finding in the visual perception literature (see Halpern, 2000 for a review) that would suggest men read microblog content differently than women.

Our user sample reported that they had been using Twitter for an average of 23 months (SD = 12.90). All

participants reported accessing their Twitter feed at least once per day on average, with 95% of them reportedly doing so multiple times per day. Forty-five percent of users reported posting tweets multiple times per day, 20% reported posting tweets once per day, and the remaining 35% reported posting tweets about once a week. The number of followers our users had varied greatly from 27 to almost 4,000. The number of users' friends on Twitter also varied greatly from 51 to over 2,000. The median number of followers was 129, and the median number of friends was 168.

### Procedure and Measures

We built a custom application that allowed each participant to log in to his or her Twitter account. The application acquired the 100 most recent tweets from the participant's feed (excluding tweets posted by the user) and used them to construct three separate tasks. Participants were tested individually, and they were asked to refrain from accessing their Twitter feed for 24 hours prior to the study in order to ensure that the tweets being shown to them were new.

#### Task 1 – Twitter Feed Reading

The most recent 50 tweets were arranged onto 10 pages of five tweets; each tweet was presented within an equally-sized rectangle that occupied approximately 5% of the total screen area. Figure 1 shows a screenshot of the region of the screen shown to participants that contained each page of tweets. Participants were shown one page of tweets at a time and were asked to read through each page as they normally would when accessing their Twitter feed, taking as much or as little time as they desired. We created our own paginated application to facilitate the collection of our eye-gaze and recognition memory metrics. In most respects, the look and feel of the experience was very similar to that of the TweetDeck application.

Participants were unable to click through linked content in the tweets or to reply or to retweet them. This restriction maintained focus on the tweets themselves rather than on linked content, which was critical to our research goal of understanding the crucial first pass of user attention over the content. As described below, in Task 3 participants later indicated the different ways they would have interacted with each tweet. We also logged the following information about the 50 tweets seen by participants: sender ID, tweet length, whether the tweet was a reply, a retweet<sup>1</sup>, and contained a link or hashtag.

As participants read each page, their eye movements were tracked using a Tobii X-50 eye tracker (50Hz sampling rate), which was calibrated for each participant before the start of Task 1. The eye-tracker was paired with a 17" LCD monitor (96 dpi), set at a resolution of

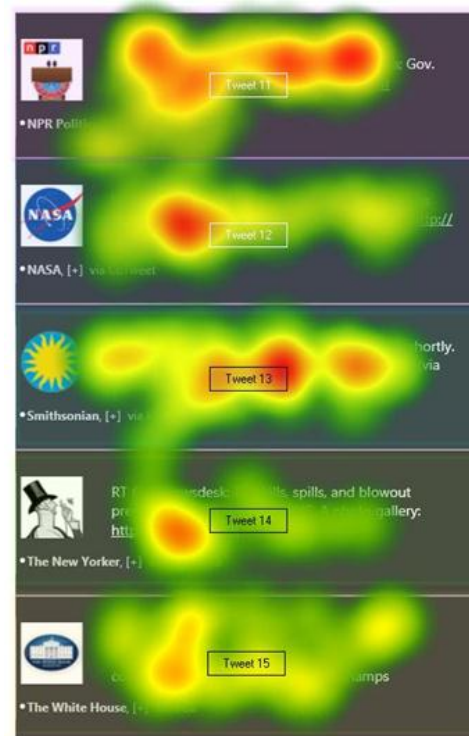


Figure 1. Example heat map visualization of time spent looking at tweets (red/darker coloring indicates longer fixations).

1024x768. Using the Tobii eye-tracker's software (Tobii Studio 1.7.3), we calculated the total number of fixations each participant made in each tweet area, the number of separate visits their eyes made to each tweet, and the total time each participant spent looking at each tweet. Fixations were defined as when the eyes focused on a 35-pixel area for at least 100ms (see Figure 1).

Of the eye-tracking measures we collected, we decided to focus on total looking time (summed fixation duration). A fixation is the most basic proxy for attention capture. Presumably, the more fixations users make, or the longer those fixations are, the more cognitive resources they are expending to encode that information, and the number of fixations was highly correlated with the total looking time for a tweet ( $r = .94, p < .001$ ). We also considered various metrics that controlled for tweet length when considering total looking time, such as dividing the number of milliseconds spent looking by the number of characters in the tweet. However, such metrics were no more informative than the overall looking time measure. Tweet length and total looking time were not strongly correlated, and such "looking time per length" metrics added to the assumptions we had to make (e.g., there should be a monotonic increase in looking time for each additional character). Thus, total tweet looking time is the only eye-tracking metric discussed here.

<sup>1</sup> Retweets were identified using a regular expression match that checked for the presence of the word "RT" or "via" coupled with the "@" character.

Measure		Result
Tweets Remembered		68.6%
Rating	High: 6-7	14.9%
	Medium: 3-5	48.1%
	Low: 1-2	36.7%
Behaviors	Click	51.7%
	Reply	3.6%
	RT	5.0%
Mean Looking Time	Total Session	2.44 min
	Per Tweet	2.92 sec

Table 1. Overview of participants' memory performance, response behaviors, and looking times.

### Task 2 – Tweet Recognition

When participants had finished reading the 50 tweets, they were then asked to perform a recognition task. The 50 tweets were intermixed with the other 50 tweets that the application had initially acquired from the participant's Twitter feed. Participants were shown all 100 tweets in random order, and for each tweet they were asked to indicate whether the tweet had been present in Task 1.

Memory for tweets is important because if information is not remembered by users it cannot be put to later use, meaning that information dissemination and influence intentions of tweet authors will have been in vain. This may be especially true for organizations and businesses who try to use Twitter to gain support. Because we measured memory with a "Yes/No" recognition task (making chance performance 50%), and the task followed immediately after participants read their Twitter feed, it should have been relatively easy. Thus, if a particular tweet wasn't remembered it is logical to assume it was not successfully encoded during Task 1. We also assume that a recall test targeting longer term memory (e.g., a day or week later) would be even less accurate.

### Task 3 – Tweet Ratings & Behavior Reporting

Lastly, after completing the tweet recognition task, participants were again shown the set of 50 tweets that they were shown in Task 1 in their initial order. Beside each tweet were questions that asked participants to rate how interesting each tweet was to them on a seven-point scale (1 represented very uninteresting, and 7 represented highly interesting), who the tweet author was (personal contact, organization, or a celebrity/other), and which behaviors they would have performed for each tweet if the interface had allowed them to (reply, retweet, click link). They

reported these behaviors via Yes/No responses, as well as a Maybe response for the reply and retweet questions.

## Results

### Attention and User Cognition (RQ1)

A description of overall results for our three measures across all types of tweets is presented in Table 1. People spent a relatively brief amount of time reading through their tweets, allocating only a couple of seconds to each tweet. Participants recalled less than 70% of what they saw, and rated about 15% of tweets as highly interesting. Figure 2 shows the relationships between these measures. As tweets were looked at longer, they were also significantly more likely to be remembered ( $p < .01$ ); likewise, tweets that were remembered were looked at longer ( $p < .01$ ). This suggests that greater attention is indeed related to deeper encoding in the context of microblogs. Also, tweets that were rated highly were looked at longer ( $p < .01$ ) and remembered significantly more often ( $p < .01$ ) than tweets that were not.<sup>2</sup>

### User Behaviors (RQ2)

Figure 3 corresponds to results in this section.

#### Clicking Links

Of the 59% of tweets containing links, participants reported wanting to click through about half of these on average (52%). Participants rated tweets containing links they would click as being significantly more interesting than tweets without links ( $p < .01$ ). There were no significant differences in looking time or memory for tweets with links that users would and would not click on. This is probably because when a tweet contains a link, that link is often shortened to a unique URL that cannot be deciphered and therefore contains no additional information that users should spend time looking at.

#### Would Reply

As reported in Table 1, participants indicated that they would reply to very few tweets overall (3.6%), with many users reporting they wouldn't reply to any tweets. Despite the miniscule number of "Yes" and "Maybe" reply responses, there were significant differences within our measures<sup>3</sup>. The group of tweets that participants would definitely reply to were remembered significantly better

<sup>2</sup> All tests are paired t-tests. For readability we report p-values only. To account for the fact that we ran multiple tests, throughout the paper we adjusted our significance criterion according to the number of tests run for each measure in each category of research question (about  $p < .006$ ). We report uncorrected p values, but refer to an effect as significant only if it meets the adjusted criterion. P values near or below the traditional .05, but above the adjusted criterion, we refer to as marginally significant.

<sup>3</sup> We decided to use paired t-tests for these comparisons rather than repeated-measures ANOVAs because some of our participants made only "Yes" and "No" responses, some made only "Maybe" and "No" response, and some made all three, so we wanted to include as much of our data as possible in our analyses. Per footnote 3, we corrected for the number of comparisons.

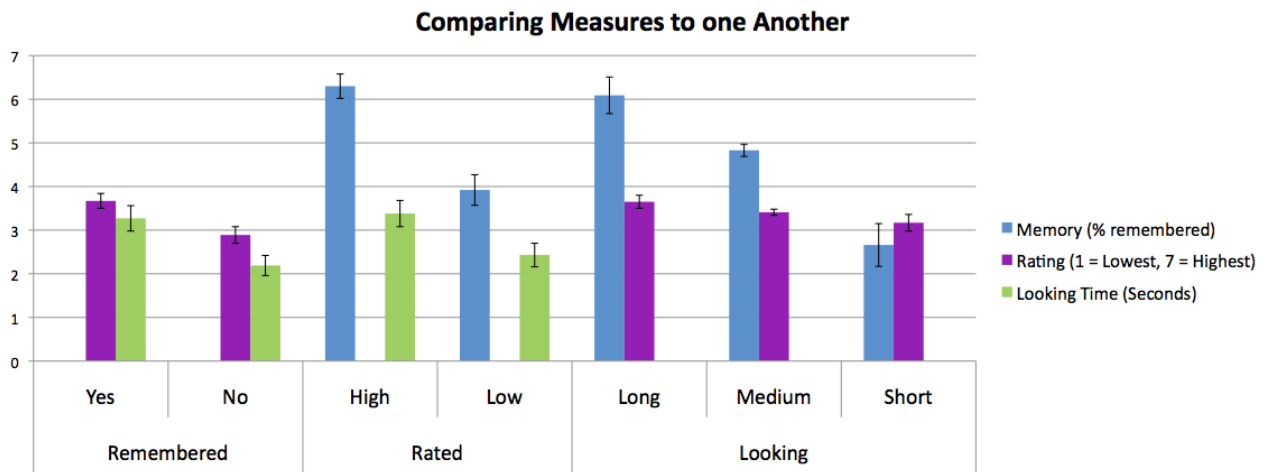


Figure 2. Each measure is plotted against the other to compare patterns across measures. Note that % remembered is scaled to be out of 7 (i.e., 7 = 100%) to facilitate comparison to the other measures.

than tweets that participants would not reply to ( $p < .01$ ), while “Maybe” and “No” were marginally significantly different ( $p = .02$ ), and there was no difference between “Yes” and “Maybe.” Also, “Yes” and “Maybe” tweets were rated as significantly more interesting than tweets that participants would not reply to (No vs. Yes:  $p < .01$ ; No vs. Maybe:  $p < .01$ ), but again the ratings for “Yes” were not significantly different from “Maybe.” Lastly, tweets that participants would reply to were looked at marginally significantly longer compared to tweets that participants would not reply to ( $p = .04$ ).

These results suggest that reply behaviors do indeed reflect interest, attention, and deep information encoding. However, the number of such behaviors actually made is extremely small, meaning that much of the high quality content encountered by participants is not reflected in their replies.

#### Would Retweet

The small percentage of tweets that participants in our study would reportedly retweet (5%) received a very high interest rating overall. When participants reported “Yes” or “Maybe” wanting to retweet, these tweets were rated significantly higher than those that users did not want to retweet (No vs. Yes:  $p < .001$ ; No vs. Maybe:  $p < .001$ ), but the “Yes” category was rated the highest at marginal significance ( $p = .02$ ). There were no significant differences for looking time and only marginally significant differences in memory (No vs. Yes:  $p = .06$ ; No vs. Maybe:  $p = .05$ ).

Though limited, these results suggest that only tweets above a relatively high threshold in terms of attention and interest are considered for retweeting. This interest threshold for retweeting seems to be higher than that for replying. Again, though, retweets only reflect a fraction of

what users attended to, considered interesting, and remembered.

#### Tweet Properties (RQ3)

Figure 4 corresponds to results in this section.

##### Contains Linked Content

As mentioned earlier, the majority of the tweets that each user saw in our study contained links (59% on average). This high percentage probably reflects our sample of active twitter users drawn from a knowledge worker population. Interestingly, on average participants looked at tweets with links for less time ( $p < .01$ ), were marginally significantly less likely to remember tweets containing links ( $p = .03$ ), and rated these tweets as less interesting than tweets not containing links ( $p < .01$ ). Thus, although link sharing via tweets is common, tweets with links do not appear to engage the user any more than tweets without links.

##### Contains Hashtag(s)

On average, 18% of the tweets each participant saw contained a hashtag, with each participant seeing 3-17 tweets with hashtags in their feed during the study. Whether or not a tweet contained a hashtag showed no reliable difference in terms of perceived interest or how long users looked it. However, participants in our study were marginally better at remembering tweets without hashtags (70% correct performance) than with hashtags (61% correct performance;  $p = .05$ ).

One explanation for this result is that hashtags make the tweet more complex and thus more difficult to encode and process the information contained in the tweet, especially given that people spent roughly the same amount of time looking at tweets with and without hashtags. Alternatively, it could be that hashtags provide users with a short-cut keyword to help them determine whether or not the tweet is interesting or relevant to them. Thus, users may give

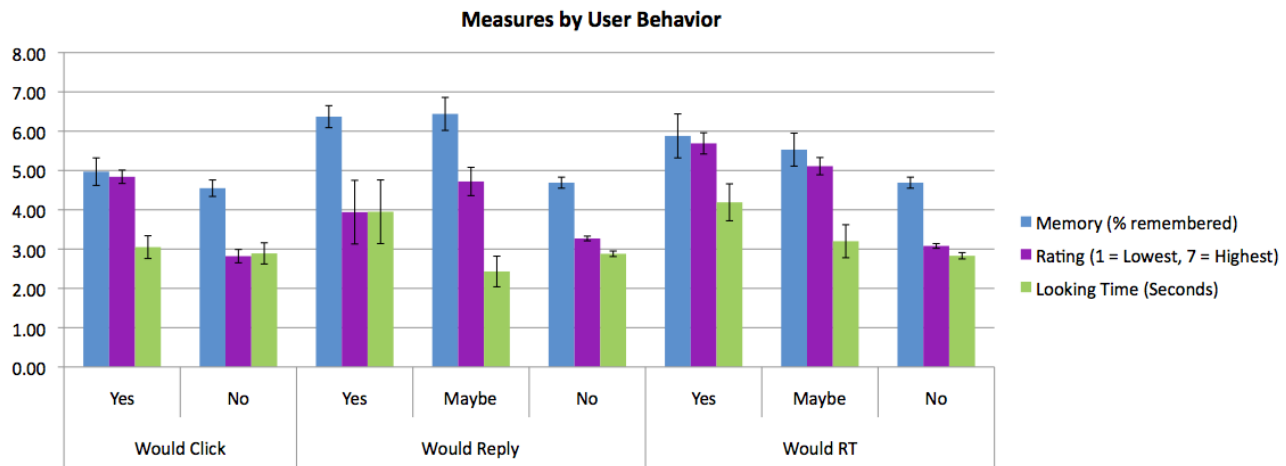


Figure 3. Mean looking time, user ratings, and memory performance are shown for each type of behavioral response that participants would reportedly make for the tweets seen during Task 1.

little attention to hashtag-containing tweets except for those few having tags that interest them.

#### Is a Reply

We found no statistically reliable differences in memory, rating, or looking times when comparing tweets that were replies to another Twitter user compared to tweets that were not replies. However, it should be noted that only eight of our participants saw a reply tweet in their stream, and the average percentage of reply tweets that these participants saw in their feeds was extremely low (2.3%), so these data are based on only a handful of items.

#### Is a Retweet

Retweets were relatively common; they made up about 20% of all tweets seen by participants in the study, and every user saw at least one retweet in their Twitter feed during the study. When a user retweets, it is probably because he or she has deemed the content of a tweet to be especially interesting (as suggested above). Interest doesn't appear to go both ways, however. Although retweets were looked at marginally significantly longer ( $p = .06$ ), people rated retweets significantly *less* interesting than "original" tweets ( $p < .01$ ). There was no significant difference in memory for original and retweets.

#### Sender Category

Across all participants the number of tweets from friends falling into different categories was relatively similar (33% personal contacts, 38% organizations, and 27% celebrities/other). There were no statistically significant differences when comparing tweets from the three friend categories for looking time or rating (e.g., tweets from celebrities were not more or less likely to be rated more highly or looked at longer). However, there were differences in memory, such that tweets from personal contacts were more likely to be remembered than tweets from organizations ( $p < .01$ ) or celebrities ( $p < .01$ ).

#### Frequency of Sender Appearance

Some authors tweet far more often than others, and there was a lot of variability among our sample in terms of how many unique authors appeared in a user's Twitter feed. Overall, it was most common for authors to appear less than five total times throughout the Twitter feed (77.2% of all tweets). For 14 of our 20 participants there were enough "frequent tweeters" in their feed for us to make a comparison between frequent authors (whose tweets appeared five or more times in feed) and infrequent authors (appeared only once). Participants looked more than a second less at tweets from frequent authors (3.74 seconds for infrequent and 2.48 seconds for frequent;  $p < .01$ ). They also remembered these tweets marginally significantly less often ( $p = .04$ ). These results give some evidence that when an author tweets frequently, each individual tweet risks being overshadowed, the result being that each of the author's tweets receives less attention.

#### Limitations

Each of our measures contained some amount of error. For memory encoding, factors that have little to do with the item's actual substance, such as the distinctiveness of an item relative to its context, can facilitate attention capture and the transfer of information from working memory to long-term memory (Baddeley, 1997). For our behavioral measures, capturing intended behavior is obviously not as precise a measure as actually recording these behaviors as they occur. However, because it was important to have a "pure" measure of looking time while participants were reading their Twitter feeds, post-task behavior reporting was the best option. Lastly, though eye-tracking unequivocally shows where participants' eyes were positioned while reading their Twitter feeds, we cannot guarantee that the looking behaviors exhibited by

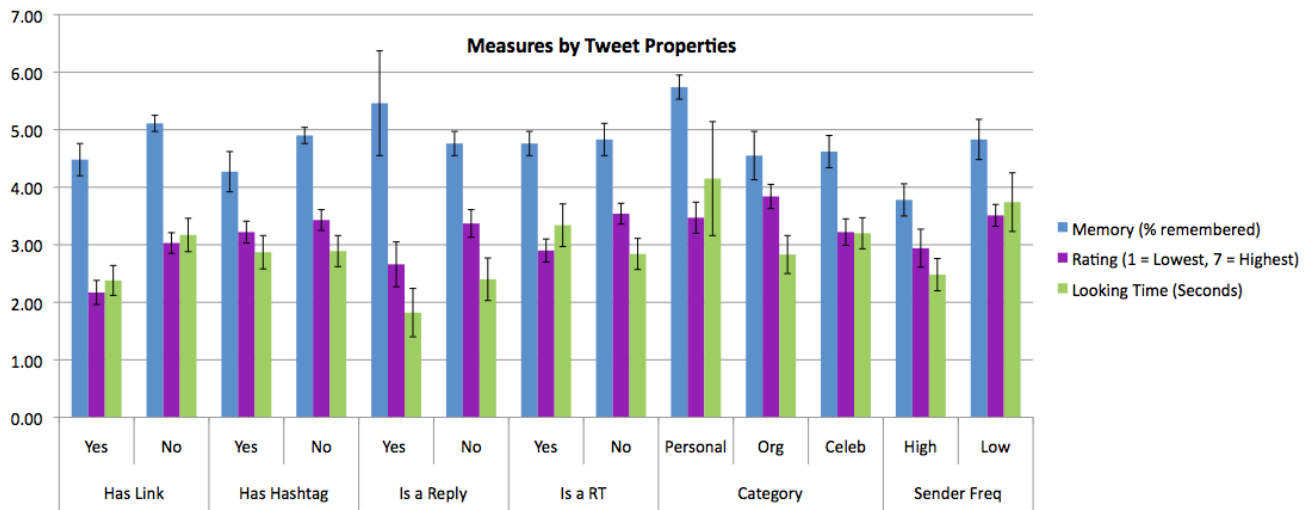


Figure 4. Total looking time, rating, and memory performance measures for tweet attributes. Note that memory performance was scaled to be percent correct out of 7 to facilitate comparison.

participants during our study were fully representative of how they might read their feed normally.

Also, a note about our sample population: as mentioned, our sample population consisted of active Twitter users who worked in a corporation. While they likely represent a “professional” type of user, they probably use Twitter differently than other groups (e.g., teens).

## Discussion

We started with the goal of analyzing three alignments: those between a user’s visual attention and 1) other cognitive measures, 2) behaviors, and 3) properties of the microblog content. Our results suggest the expected relationships, but also expose gaps that highlight difficulties in the basics of microblog consumption and may be relevant to network level analyses of Twitter.

First, generally speaking, people tend to spend only a brief amount of time (3s) reading each tweet in their Twitter feed, little content (15%) is considered highly interesting, and much of the content users encounter when reading their feed is forgotten a few minutes later. Further, factors like the type of author impact memory for content (tweets from friends were remembered better than those from organizations). This result bears on information contagion in that the effective reach of information may be constrained by a fairly weak human memory for microblog content: Even content seen by users may not be remembered and thus may have little effect.

Second, though replies and retweets do appear to reflect user interest and attention, users would reply to (24%) and retweet (34%) only a fraction of the tweets that they rate as highly interesting, and a much smaller percentage of tweets rated at medium interest. This means that much content considered interesting is not being captured by measures of

user behavior. The relationship between these percentages and network level measures that bear on topic interest (e.g., trending topic identification, information diffusion) is an area for future exploration. For instance, does this interest that is not being captured “average out” across users such that the topics of greatest interest overall are still surfacing? If not, are there individual differences that might be leveraged, such as giving additional weight to retweets from users who rarely retweet?

Our results also expose a number of misalignments in the current microblogging environment. Ideally, these gaps can be minimized by bringing to the user’s attention the content of most value. First, even though users tend to only retweet content that they consider interesting to them, they don’t find content retweeted by their friends to be more interesting. They do, however, look marginally longer at retweets. Thus retweets are commanding an amount of attention disproportionate to their level of interest, suggesting that a simple filter for retweets versus original content could be beneficial. Second, frequent tweeting by an author decreases, rather than increases, the amount of attention any individual tweet of the author receives. In addition to suggesting a bit of caution when tweeting heavily, this also suggests that microblog clients should highlight infrequent authors.

Third, despite the fact that Twitter has become a prime medium for link sharing, our results suggest that people tend to pay less attention to tweets with links, as evidenced by the shorter amount of time people spent looking at them and the overall less interest in and poorer memory for tweets with links. However, participants did report wanting to click on about half of the links that they saw, and they rated that those tweets higher. Thus some tweets with links command attention, but the high percentage implies user difficulty in evaluating how interesting a tweet is if its

primary message can only be comprehended by clicking through to linked content. For system designers this suggests a need for link previewing, either in-line with tweet content or in response to user action.

Finally, hashtags generally are used to provide a category or label for a tweet, and one might predict that hashtags would thus support memory for tweets. However, our results show that tweets with hashtags were remembered marginally less well than tweets without hashtags, suggesting that users may not encode hashtags as well as standard tweet content. This is an area for additional exploration, and this finding should be reconciled with Suh et al. (2010) who found that the presence of hashtags increases the likelihood of a tweet getting retweeted. Perhaps people are more likely to remember the content they retweet, which often contains hashtags, but the hashtags of others do not aid memory.

Taken together these findings point to a disconnect between the diffusion and presentation of information and the way users are attending to that information. With respect to diffusion, the primary vehicles for user-based information diffusion—retweeting, frequency of tweeting, and link sharing—all showed counterproductive results: retweets were not seen as more interesting, frequent tweeting reduced per tweet visual attention to the author, and including links in tweets decreased visual attention to the tweet on average. Therefore, from a content presentation standpoint, equality in tweet presentation may be suboptimal, as even completely uninteresting content garners some user attention. To address these issues, filters and interface elements that help users decide what content is most interesting may be useful (see Hong et al., 2010 for somewhat related examples).

## Conclusion

Microblogging is a dynamic environment for sharing information, but it presents unique cognitive challenges for users. This work focused on the critical moment when users attend to pieces of information in their microblog stream: what do they attend to and how does that relate to what they find interesting, what they encode in memory, and what they actions they intend to take. We found that users only spend about three seconds reading each tweet, and generally speaking they are able to use this three seconds to attend to content they find interesting and remember. However, inefficiencies remain: a large percent of interesting content that is forgotten, user behaviors (replies, retweets) reflect only a fraction of highly interesting content, and properties of the content itself often decrease attention rather than help users quickly direct their attention to the content of highest value.

We see two directions for utilizing these results. First are the various ways they might be integrated into network scale analyses. At a minimum, these results point to noise (e.g., interest not perfectly related to behavior) that can be incorporated as error terms, as well suggest variables to incorporate into models (factoring in the type of sender given its impact on memory). The second direction is to target interface designs that can help reduce the noise. Here we suggest designs that focus on expanding information depth, with the goal of helping users make accurate decisions about interest level of tweet content in a short amount of time.

## Acknowledgements

We thank Aaron Hoff for building the study application.

## References

- Baddeley, A.D. (1997). *Human memory: Theory and practice*. Psych. Press: Hove, UK.
- Bernstein, M. S., Suh, B., Hong, L., Chen, J., Kairam, S., & Chi, E. H. (2010). Eddi: Interactive Topic-Based Browsing of Social Status Streams. *Proc. UIST 2010*.
- boyd, d., Golder, S., & Lotan, G. (2010). Tweet, Tweet, Retweet: Conversational Aspects of Retweeting on Twitter. *Proc. HICSS-43*
- Buscher, G., Cutrell, E., & Morris, M. R. (2009). What do you see when you're surfing? Using eye tracking to predict salient regions of web pages. *Proc. CHI 2009*.
- Chi, E. H., Gumbrecht, M., & Hont, L. (2007). Visual foraging of highlighted text: An eye-tracking study. *Proc. HCI 2007*, 589-598.
- Craik, F.I.M., & Lockhart, R.S. (1972). Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior*, 11, 671-684.
- Cutrell, E., & Guan, Z. (2007). What are you looking for? An eye-tracking study of information usage in web search. *Proc. CHI 2007*.
- Halpern, D.F. (2000). *Sex Differences in Cognitive Abilities*. Psych. Press: Hove, UK.
- Hong, L., Convertino, G., Suh, B., Chi, E., Kairam, S. (2010). FeedWinnower: Layering Structures Over Collections of Information Streams. *Proc. CHI 2010*.
- Just, M.A., & Carpenter, P.A. (1976). Eye fixations and cognitive processes. *Cognitive Psychology*, 8, 441-480.
- Kwak, H., Lee, C., Park, H., & Moon, S. (2010). What is Twitter, a Social Network or News Media? *Proc. WWW 2010*.
- Suh, B., Hong, L., Pirolli, P., & Chi, E. (2010). Want to be Retweeted? Large Scale Analytics on Factors Impacting Retweet in Twitter Network. *Proc. SocialCom 2010*.
- Yang, J., & Counts, S. (2010). Comparing Information Diffusion Structure in Weblogs and Microblogs. *Proc. ICWSM 2010*.
- Yang, J., & Counts, S. (2010). Predicting the Speed, Scale, and Range of Information Diffusion in Twitter. *Proc. ICWSM 2010*.