

# Is Twitter a Good Enough Social Sensor for Sports TV?

Venu Vasudevan, Jehan Wickramasuriya

Betaworks, Applied Research Center

Motorola Mobility, Libertyville IL

{venu.vasudevan, jehan}@motorola.com

Siqi Zhao, Lin Zhong

Dept. of Electrical & Computer Engineering

Rice University, Houston TX

{siqi.zhao, lzhong}@rice.edu

**Abstract**— The globalization of TV programming (especially in the Sports and Reality segments) is leading to a bewildering amount of choice for TV watchers. The social currency around knowing what's happening in these programs as they happen combined with notable undulations in the interestingness of these programs leads to a navigation problem for the TV watcher. The pervasive use of Twitter in conjunction with TV watching makes it a potential sensor for real-time TV, and therefore a building block in solving the user problem of interstitially navigating to the most interesting program at any point in time. Given that users have to tune into shows with peak moments as soon as those peak moments happen, interstitial navigation has to be high enough performance to enable TV event detection and user tune in within tens of seconds of the event.

The question being addressed in this paper is – *Do Twitter's social sensing capabilities have sufficient precision and timeliness to cater to the 'extreme' navigation needs of sports fans?* And if so – how can we design a TV event detection framework that can be extended to multiple sports, and beyond sports. We focus on navigation for sports programming in the narrative, as an extremely demanding TV application that also has high market attractiveness. However, we anticipate that the ideas and architecture proposed herein apply to any TV programming that lends itself to interstitial viewing, and elicits a high level of real-time user participation in social networks.

**Keywords:** *social sensors, video navigation, collaboration, context, crowd-sourcing*

## I. INTRODUCTION

The quantity and quality of mass human participation in social networks (particularly Twitter) has led to Social Networks emerging as a real-time, human sensor network (HSN) for the physical world [1,2]. This ability of decentralized humans to rapidly sense even weak signals in the physical world and share this information in timely ways on Social Networks such as Twitter has been transformational to a number of disciplines. For instance, it has transformed Journalism & Seismology [13], fields that rely on the reporting of physical events in an accurate and timely manner. Digitally connected humans now serve as a force amplifier for journalists who are limited in their ability to be in the right place and the right time, but can now use

the HSN as both a sensor and actuator to accelerate the speed and scope of a news story.

The scope of Twitter user behavior and consequently that of social sensing has expanded from real-time, in-venue sensing of *physical* events to real-time user response to *digital* events such as TV shows. In cases where the TV broadcast is near real-time (e.g. a football game, a reality show, or widely followed events such as the Oscars), the source of tweets for a physical event now encompasses not only the onsite audience, but also a larger and geographically dispersed audience of TV viewers who are digitally co-present with the in-venue audience. A driver for such 'TV tweets' is the pervasiveness of *dual-screen* TV viewing [14], wherein viewers watch TV while simultaneously interacting with social networks (notably Twitter) on the second screen.

This emerging behavior around TV Tweets has the potential to address a growing problem for TV users – that of *interstitial* (just-in-time) navigation for TV content based on the program's *interestingness* (i.e. 'where the action is' at this moment). Interstitial TV watching is trending up because some of the more popular TV genres such as Sports and Reality TV are *atomized* (either by design or in emergent fashion) for viewing in segments rather than from beginning to end. A sporting event for instance is atomized in unplanned fashion into players and plays, whereas a reality show such as *American Idol* is pre-orchestrated to create interesting moments around specific singers, mentors or judges. Taking Sports TV as the example – a user who is juggling multiple concurrent games on TV is faced with two unsatisfactory choices – that of randomly switching between televised games, or that of recording one or more of these games for deferred viewing. The social nature of sports makes the latter ineffective, as deferred viewing of a game (whose result you don't want to know in advance) requires disconnecting from your social network for a period of time lest one of your friends informs you of the result by accident. Thus there is a need for an intelligent program guide that enables a user to tune into whichever sport is most interesting at the moment.

Further, *interestingness* of a game is a nuanced concept that goes beyond **strong signals** [10] – major and infrequent

in-game events such as a touchdown in American Football (or a goal in Soccer). **Weak signals** such as a spectacular *sack* in American Football or a *give-and-go* in Soccer are of high interest to fans even though these events don't contribute to the final score. In fact, weak signals are the majority of a sport's lexicon in both quantity and occurrence, and a key part of what draws and sustains fan engagement. A significant market force driving user interest in weak signals (especially in sports) is the phenomenon of fantasy sports.

Fantasy sport is a game overlay on any sport (e.g. American Football, Basketball etc.) where fans pick and manage their personal virtual team (also known as a Fantasy team) during the season. Because a fan's fantasy team can comprise of players from different actual teams, fan allegiance shifts from singular games between their favorite teams to any games that feature the players in their fantasy team. Furthermore, fantasy points are often associated with weak signals (as can be seen in Fantasy Football scoring sheets). Anecdotal information on the user base and market size of Fantasy Football alone estimates that about 30 million active Fantasy users that spend about \$1.5B directly on Fantasy related products. The combination of an increased population of sports events and the rise of fantasy sports thus makes a strong case for the need for a real-time social navigation experience for TV for users. Specialized channels such as the NFL Redzone have emerged as a curated channel, where manual curators watching multiple games switch you to one or a few of the most interesting moments in ongoing NFL games. This experience has been incredibly popular, having debuted in 2005 via DirectTV and then adopted by the NFL Network in 2009 and continues to grow in viewership with the prominence of fantasy football. While NFL Redzone is in the spirit of what we are trying to achieve, it does not personalize the user experience to a particular fan's teams or players, which is something central to our goals around personalized, interstitial TV and most notably - manually curated which is difficult to replicate for multiple sports and genres.

The key question being addressed in this paper is - **Can we create an effective social navigation experience for TV via real-time analysis of data from social networks (such as Twitter)?** This question can be broken down into the following components (with each subsequent question assuming a satisfactory answer to the previous one):

1. Can we detect all (or a large portion) of the event lexicon (both strong and weak signals) of a sport with the speed and accuracy required for a real-time TV experience from the Twitter HSN
2. How do we efficiently identify the most capable Human Sensors for a particular event in the event lexicon (e.g. *sack*. vs. *blitz* in American Football)
3. How do we reduce the total effort in aggregating the optimal set of Human Sensors from an HSN that cover the event lexicon for the program
4. What is the quality of sensing (per event and in aggregate) that we can expect from the HSN

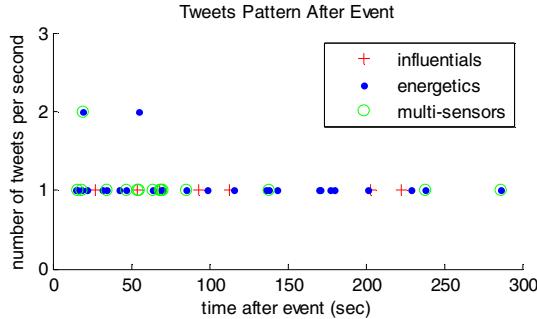
## II. RELATED WORK

Related work falls into two broad categories around socialization and information diffusion. The first category comprises of work in Social TV that investigates the patterns of socialization associated with real-time TV watching. Work in this category pertains to TV and is somewhat agnostic of the social network. The second category comprises of work that studies the dynamics of HSNs (Twitter-specific), and the identification of effective human sensors in efficient ways.

In the first category, work in [6] makes the point that real-time and deferred interaction with video content is a rich vocabulary, and that conversation is inherent to mass media content. Case studies provided in [7] and [8] look at specific forms of real-time social interactivity around TV, and ways in which these social sensor readings are distributed across the TV community via both public and proprietary social networks. Work described by [9-12] makes the case that Twitter is a particularly promising sensor for TV, and that tightly interleaved TV watching and tweeting cuts across genres, but is particularly intense in sports. Collectively, the work in this category provides support for TV Navigation as a problem and HSNs as a solution. In particular, this work illustrates the efficacy of the real-time detection of *strong events* using HSNs.

In the second category, work in [3-5] identifies a credible, topic-specific HS quorum that is both accurate and timely in its detection of the topic. If one views TV sub-events as topics, our goal of finding effective human sensors of TV sub-events is closely related to understanding the information detection and diffusion patterns in a HSN. However, TV sub-events have two distinct properties of being *objective* and *super-ephemeral*. They are *objective* in the sense that one doesn't need deep expertise to recognize TV sub-events, and there is no social or commercial value in lying about them, thus eliminating the false positives that are deliberate. The tribal nature of sports does lead to *conditionality* and *asymmetry* of sensing (e.g. I'm more likely to Tweet about the clever moves of my team, rather than the opponent, or the failures of a player I dislike rather than a player I'm neutral about). This means that human sensors might be reliable for a topic (e.g. fumbles in American football) but one-sided (i.e. report a fumble only when it happens to the opposing team).

Where the objective nature of TV sub-events makes them easier to detect than events where expert selection is a bigger problem, their *super-ephemerality* leads to a unique challenge. By *super-ephemerality*, we mean the extreme latency constraints for human sensing, even compared to real-world 'real-time' events (such as earthquakes). Unlike physical world events where detection is sufficient in minutes or hours, TV sub-events need to be detected with a latency ranging from seconds to tens of seconds, to be a tool for Social Navigation. In the next section we discover that the objective and super-ephemeral nature of TV sensors



**Figure 1** Sample Tweet pattern after a sack from an NFL Game: Chicago Bears vs. Green Bay Packers.

leads us to measure human sensor effectiveness somewhat differently from the majority of related work on human sensing via HSNs.

### III. SYSTEM & EXPERIMENT

In this section, we define our notion of an effective TV sensor, and describe the experimental setup that we put in place to test the hypothesis across multiple professional football games associated with the NFL (National Football League). To test the generality of the hypothesis, we tested its efficacy in event detection across head, shoulder and tail terms in the Zipf distribution, based on a representative sample of football vocabulary taken from the American Football Glossary [17].

In terms of a sports vocabulary, *head* terms are those that typically appear on the scoreboard (e.g. touchdown and field goal) and are therefore widely understood and are sensed by a very large population of human sensors. A touchdown in a NFL game typically elicits several hundred thousand tweets, as can be seen from some of the live experiments captured at <http://sportsense.us> - a website we created that provides real-time sensing of sporting events using data obtained from Twitter. A heuristic we use to determine shoulder terms is to regard any term that is given a score in Fantasy Football leagues (e.g. Fumble, Sack) as shoulder terms. *Shoulder* terms are both well known and well sensed because they are part of the vernacular of football fans, and further influence a fan's Fantasy Football score. Shoulder terms can typically (but not exclusively) be found in Fantasy Football scoring sheets. *Tail* terms (e.g. Chop block, Wildcat formation) are football phenomena that are interesting enough to cause a football fan to change channels, but also nuanced enough that only a very small population of football mavens will both recognize and sense these events in a timely manner.

#### A. What makes an effective TV Sensor?

In traditional information retrieval (IR) terms, human sensors with both high *precision* (low false positives) and high *recall* (consistency of sensing) are desirable. Additionally, given the super-ephemeral nature of the TV navigation experience, we'd like low latency and ideally

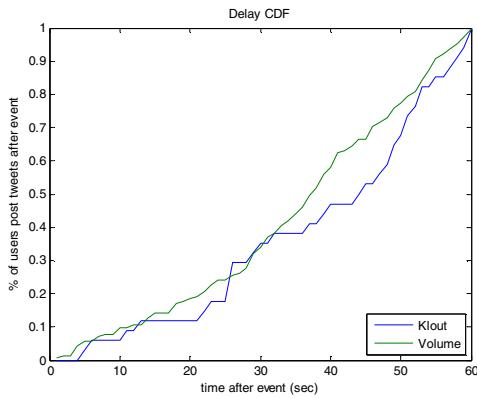
perfect *quantized latency*. By quantized latency, we mean a latency measure (in seconds) converted into an enumerated type to depict the extent to which the user (who navigates to the associated TV channel) can recapture the moment. A simple form of quantized latency we use here is to convert any latency metric to either 0 or 1 depending on whether its value is less than or more than 1 minute (and in some cases, 30 seconds). This captures the intuition that a number of cable TV channels in the US allow a user to rewind the TV channel inline by 30 to 60 seconds, so a quantized latency value of 0 enables the user to replay the event without loss of information.

For the purposes of this work, we divide the Twitter human sensor population into two classes – **influentials** and **energetics**. *Influentials* are users with perceived expertise on the topic (in this case a football event), with expertise defined by the Twitter population. *Energetics* are users with a high propensity to tweet, as indicated by past tweeting behavior as it relates to football in general, are well as term-specific tweeting behavior from past football games. The intuition is that energetics correlate positively with recall, where influencers correlate positively with precision.

While measurements of influence in Twitter are an evolving space [16], we pick **Klout** [18] scores as the measure of influence due to its status as a de facto standard on Twitter. Klout scores of Twitter users are widely available, and while Klout has some well known and highly publicized drawbacks [20], it has served well as a first approximation of the reach and impact of user Tweets on their listening population. We should note that our analysis is not dependent on the specific reputation metric (i.e. Klout) and can utilize alternatives if necessary.

For energetics, we attempt two kinds of scoping functions to filter high volume Tweeters in sports – *all-terms* and *each-term*. In the *all-terms* model, the energetics are filtered based on the total tweet volume across all the football terms that are part of the experiment. In the *each-term* model, the energetics are filtered based on their tweet volume for that term relative to the population. Since Klout is a single aggregate score and doesn't support a measure of term-specific influence, the influentials are currently the same in both these models. However, a part of future work is to consider a derived Klout metric that accounts for a tweeter's conditional influence on a particular term.

Lastly, we calibrate human sensors based on sensor diversity into multi-sensors and uni-sensors. Multi-sensors make it to the list (either influentials or energetics) on multiple terms and therefore are able to detect multiple events with high performance. Uni-sensors are specialized to one kind of event. We'd like to understand the relative performance of multi-sensors, and whether observing more things makes a human sensor overloaded (and therefore slower) or more committed to timely reporting (therefore faster).



**Figure 2 Delay CDF shows that energetics are faster than influentials within 60 seconds after events.**

### B. Data Collection Setup

We developed approaches to effectively gather Twitter data, due to the fact that we are measuring weak signals and rate-limited by the Twitter public API. In comparison to strong signals (e.g. a touchdown in the Super Bowl or a new year tweet) that generate a spike of several million tweets, weak signals are somewhat harder to clearly distinguish in terms of tweet volume. The weak signal terms include shoulder terms such as interception, fumble, sack, and tail terms such as onside, wildcat, roughing, and blitz. To track the weak signals, we fed these weak signal terms into the Twitter Streaming API, and generated a pattern-matching lexicon using those weak signal terms. In the process, we collected about 3 million of tweets from 15 prime time games in the first 7 weeks of 2012-13 NFL season.

Among the Twitter APIs (Search, REST, Streaming), only the Streaming API provides a continuous flow of tweets in real-time and, therefore is appropriate for our experiments. Streaming API provides three methods, i.e. sample, firehose and filter. Sample method returns a small random sample of all public statuses. Firehose method returns all public statuses but is not publicly available. The filtered method of the Streaming API returns public tweets that match one or more filter predicates, including user ID, keyword or location. The *track filter* predicate allows us to collect tweets using keywords, which is critical to our lexicon-based methods, and ensures that all tweets returned are the latest available tweets.

Two interesting weak signal related issues we encountered (and addressed) along the way pertained to lexical conflicts with non-football terms, and dealing with the limitations of Twitter's public API (parameters and rate limit). On the former, our weak signal lexicon included words (e.g. blitz and sack) that were also common English words with popular uses unrelated to their football usage. However, we found that limiting the tweet filtering to game times and adding a small amount of context to the term extraction was adequate to weed out most non-Football uses

of these terms. The public Twitter Streaming API has two major types of limits, parameters and rate limit. The parameters are limited to up to 400 track keywords, 5,000 follow userids, and 25 location boxes. While the rate limit imposed by Twitter's Streaming API of 50 tweets per second may not seem like much, it was more than adequate given the bounded signal vocabulary along with the weakness of the signals we were tracking.

## IV. RESULTS

In this section, we investigate the delay, recall and precision of influentials, energetics and multi-sensors. Our goal is to characterize these users for developing methods that can recognize weak events and enable social navigation experiences for TV. Since weak signals are so weak (in terms of Tweet volume), conventional methods [10] do not work for two reasons. Firstly, if we aggregate more tweets over a longer period of time to recognize event peaks, it considerably increases the delay. This will not meet our goal of <1-minute delay for social navigation purposes. Secondly, as the signal becomes weak, the signal to noise ratio becomes lower and the peak becomes obscure. As a result, it is hard to recognize events accurately.

As we described in the previous section, we focus on users with particular characteristics. We select the influential and energetic users from a pool of general users. Among the influentials and energetics, we further investigate sensor diversity via multi-sensors and uni-sensors. Figure 1 shows a sample tweet pattern after a sack in the game between the Chicago Bears and Green Bay Packers. For the three types of sensor characterizations we examine delay, precision and recall in detail.

### A. Energetics are 'Quicker on the draw'

We measured the event detection latency of the two classes of users on a labeled set of football events occurring in a set of NFL games (in real-time). This data was then annotated with the *pseudo-actual* times of the events of interest (contained in the football lexicon). We use the term *pseudo-actual* because there are delays in TV broadcasts (ranging from 2 to 7 seconds depending on the nature of your feed) compared to the actual occurrence of events. However, given that the preponderance of sport tweets come from a TV audience, the *pseudo-actual* can be considered the actual (zero latency) moment for the purposes of this discussion.

We extracted tweets posted by influentials and energetics within 1 minute after the occurrence of events using all-term and each-term model respectively. It should be noted that the 1-minute (and 30 second) time limit is derived based on the needs of TV navigation. Energetics post about twice as many tweets as influentials within 1 minute (after events). In both models, 10% of the influentials and 7% of energetics post tweets in this 1-minute window. Significantly, the median energetic has a latency of 33 seconds, which is 12 seconds faster than the 45 second latency of the median influential.

**Table 1 Precision and recall for different users**

	Influential	Energetic	Influential w/ multi-sensor seed	Energetic w/ multi-sensor seed
Precision	73%	61%	76%	74%
Recall	72%	77%	72%	70%
F score	0.73	0.68	0.74	0.72

Since ideally we'd like the TV navigation experience to have latencies lower than 30 seconds, the fact that the median energetic is fairly close to that is encouraging. Figure 2 shows the CDF of delay in this 1-minute window for both classes of users.

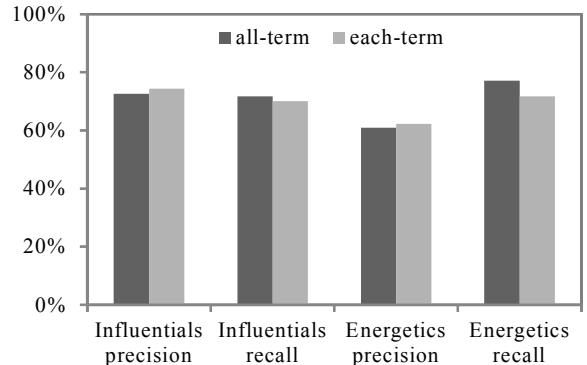
An interesting side note is the proportionally higher level of interest among influentials for tail terms. We find that more than 90% of users who tweet about “sack” are energetics but the percentage drops to 50% for “onside kick”. This observation quantifies the intuition that influentials are aficionados of the game who pay more attention to unusual occurrences.

#### B. Influentials are (slightly) better Sensors

We measured the relative quality of sensing between influentials and energetics using traditional IR vocabulary (precision, recall, and a non-weighted F-Score that gives equal weights to precision and recall).

While we maintained the spirit of the IR terms (precision indicating fewer false positives, and recall fewer false negatives), it is important to note that we customized the definitions of these IR terms to the context of solving an event detection problem with a latency constraint. We assume that event detectors are threshold functions – and event is ‘detected’ if the number of event-related tweets within a 1-minute sliding window exceeds the pre-defined threshold. The threshold itself is event-specific and derived from a training set of games. So an event detected by the threshold function between minutes T and T+1 is considered a false positive if the actual event happened at say T-5, even though an actual event happened (but outside the tolerance window). Similarly, from a recall perspective, an event that happened at T, but wasn't detected in the window between T and T+1 is considered a false negative, even if it can be detected between minutes T+2 and T+3. This is necessary due to the low volume of tweets due to constraints of the Twitter public Streaming API and the weak terminology of interest.

Table 1 summarizes the precision, recall and F-score numbers based on a corpus of 110 events across 11 games, using the event-specific threshold-based approach alluded to above. The table shows that influentials have notably higher precision (77% vs. 61%), and marginally poorer recall (72% vs. 77%) over a 1-minute sliding window. Given the CDF in



**Figure 3 Each-term model does not improve results**

Figure 2, we anticipate that the ‘recall gap’ will be larger in a 30-second window, with the recall of energetics being more than 5% better than influentials. The rationale for this intuition is the fact that 30 seconds is fairly close to the median energetic, but substantially short of the median influential. In particular, we note that over half of the influentials are official team accounts and commentators. Therefore, while their reports are of high quality, they have delays that can be as high as 3-5 minutes.

#### C. Multi-sensors Achieve Higher Precision as Seeds

Among the influentials and energetics, we also considered those users who tweet about multiple events in the game. In particular, rather than use multi-sensor as a completely separate strategy, we added multi-sensors as a *hybrid strategy*. In other words, where sections A and B considered pure sets of energetics and influentials respectively, we now impose the constraint that both these sets contain at least one multi-sensor. The intuition here is that multi-sensors are nuanced and dedicated observers of the game, and would therefore improve both precision and recall.

The rightmost columns of Table 1 show the impact of adding multi-sensors as a seed to both pure influential and pure energetic strategy. In both cases, we see an improved F-score of a seeded strategy over the pure strategy counterpart. However, there is a slight drop-off in recall between the seeded and pure strategies involving energetics. So, while the preliminary results (improved F-score with seeding) validate the intuition around multi-sensors, we need to collect more data to confirm this hypothesis more definitively.

#### D. Each-term vs All-term. More ≠ Better

In Section III A, we hypothesized that choosing the human sensors on a term-specific basis (i.e. the each-term model) might improve sensing effectiveness, as different people might be effective sensors for different terms. In practice, this does not seem to be the case. As Figure 3 shows, the each-term model seems to perform no better than the term agnostic approach (i.e. All-terms model). Therefore,

the added complexity of compiling a new set of human sensors for each football term does not yield any enhancements in human sensing performance.

## V. DISCUSSION, CONCLUSIONS & FUTURE WORK

In this paper we aimed to quantitatively and qualitatively understand if people's Twitter behavior had sufficient accuracy, timeliness and fidelity to power a navigation experience for the more 'extreme' genres of TV (such as Sports and Reality). Our conclusion so far is a qualified 'yes'.

We have included a brief overview of the data from a corpus of prime time NFL games to provide numbers behind our cautious optimism. Several of our findings validate and quantify intuition. This includes the discovery that energetics are quicker, but influentials are more accurate and nuanced. We were somewhat surprised by the 'fat-tailed' nature of the tweet distribution around football events, and that tweets around seemingly ephemeral events stretched to several minutes. Also somewhat surprising was the relative ineffectiveness of term-specific filtering of human sensors when compared to a term-agnostic approach. While both findings have partial explanations (around institutional tweets and the nature of fan engagement respectively), the explanations are incomplete and require further validation.

In ongoing and future work, we anticipate that our focus will be on increasing the precision of data collection and analysis via more nuanced lexical patterns and closer examination of weak signal extraction techniques. Other areas of near term focus are the inclusion of a larger proportion of Fantasy-related terms, system integration with broadcast TV and associated second-screen experience (with accompanying user study), and a broadening to other genres (other sports or differing genres) to understand the scope and generality of these conclusions.

## REFERENCES

- [1] Vivek K. Singh, Mingyan Gao, and Ramesh Jain. 2010. Social pixels: genesis and evaluation. In *Proceedings of the international conference on Multimedia* (MM '10). ACM, New York, NY, USA, 481-490.
- [2] Takeshi Sakaki, Makoto Okazaki, and Yutaka Matsuo. 2010. Earthquake shakes Twitter users: real-time event detection by social sensors. In *Proceedings of the 19th international conference on World wide web* (WWW '10). ACM, New York, NY, USA, 851-860.
- [3] Daniel M. Romero, Brendan Meeder, and Jon Kleinberg. 2011. Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter. In *Proceedings of the 20th international conference on World wide web* (WWW '11). ACM, New York, NY, USA, 695-704.
- [4] Byungkyu Kang, John O'Donovan, and Tobias Höllerer. 2012. Modeling topic specific credibility on twitter. In *Proceedings of the 2012 ACM international conference on Intelligent User Interfaces* (IUI '12). ACM, New York, NY, USA, 179-188.
- [5] Shaozhi Ye and S. Felix Wu. 2010. Measuring message propagation and social influence on Twitter.com. In *Proceedings of the Second international conference on Social informatics* (SoInInfo'10), Leonard Bolc, Marek Makowski, and Adam Wierzbicki (Eds.). Springer-Verlag, Berlin, Heidelberg, 216-231.
- [6] David A. Shamma, Lyndon Kennedy, and Elizabeth F. Churchill. 2012. Watching and talking: media content as social nexus. In *Proceedings of the 2nd ACM International Conference on Multimedia Retrieval* (ICMR '12). ACM, New York, NY, USA, , Article 12 , 8 pages.
- [7] N. Narasimhan, T. Horozov, J. Wodka, J. Wickramasuriya, and V. Vasudevan. 2009. TV clips: using social bookmarking for content discovery in a fragmented TV ecosystem. In *Proceedings of the 8th International Conference on Mobile and Ubiquitous Multimedia* (MUM '09). ACM, New York, NY, USA, , Article 13 , 8 pages.
- [8] Keith Mitchell, Andrew Jones, Johnathan Ishmael, and Nicholas J. P. Race. 2011. Social TV: The impact of social awareness on content navigation within IPTV systems. *Comput. Entertain.* 9, 3, Article 19 (November 2011), 29 pages.
- [9] Shoko Wakamiya, Ryong Lee, and Kazutoshi Sumiya. 2011. Crowd-powered TV viewing rates: measuring relevancy between tweets and TV programs. In DASFAA'11, Jianliang Xu, Ge Yu, Shuigeng Zhou, and Rainer Unland (Eds.). Springer-Verlag, Berlin, Heidelberg, 390-401.
- [10] Zhao, S., Zhong, L., Wickramasuriya J., and Vasudevan V. Humans as real-time sensors of social and physical events: A case study of twitter and sports games. Arxiv Preprint, 2011, [http://arxiv.org/abs/1106.4300\[11\]](http://arxiv.org/abs/1106.4300[11])
- [11] Arkaitz Zubiaga, Damiano Spina, Enrique Amigó, and Julio Gonzalo. 2012. Towards real-time summarization of scheduled events from twitter streams. In *Proceedings of the 23rd ACM conference on Hypertext and social media* (HT '12). ACM, New York, NY, USA, 319-320.
- [12] Isabel Anger and Christian Kittl. 2011. Measuring influence on Twitter. In *Proceedings of the 11th International Conference on Knowledge Management and Knowledge Technologies* (i-KNOW '11), Stefanie Lindstaedt and Michael Granitzer (Eds.). ACM, New York, NY, USA, , Article 31 , 4 pages.
- [13] Zack Whittaker. 2011. Why Twitter inherently reports news before traditional media. In ZDNet <http://zd.net/nZbk0u>
- [14] Google Multiscreen Report, Navigating the new multi-screen world: Insights show how consumers use different devices together, <http://bit.ly/O3nU6A>
- [15] Nitya Narasimhan and Venu Vasudevan. 2012. Descrambling the social TV echo chamber. In MCSS '12. ACM, New York, NY, USA, 33-38.
- [16] Anger, Isabel, and Christian Kittl. "Measuring influence on twitter." *Proceedings of the 11th International Conference on Knowledge Management and Knowledge Technologies*. ACM, 2011.
- [17] American Football Glossary of Terms. [http://en.wikipedia.org/wiki/Glossary\\_of\\_American\\_football](http://en.wikipedia.org/wiki/Glossary_of_American_football)
- [18] Klout (definition) <http://en.wikipedia.org/wiki/Klout>
- [19] Professor Sparks Controversy for Klout-Based Grading <http://bit.ly/OSKm13>