The Wisdom of Social Multimedia: Using Flickr For Prediction and Forecast

Xin Jin
Department of Computer Science
University of Illinois at Urbana-Champaign
xinjin3@illinois.edu

Andrew Gallagher
Kodak Research Laboratories
Eastman Kodak Company
andrew.gallagher@kodak.com

Liangliang Cao
Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign
cao4@illinois.edu

Jiebo Luo
Kodak Research Laboratories
Eastman Kodak Company
jiebo.luo@kodak.com

Jiawei Han
Dept. of Computer Science
UIUC
hanj@illinois.edu

ABSTRACT

Social multimedia hosting and sharing websites, such as Flickr, Facebook, Youtube, Picasa, ImageShack and Photobucket, are increasingly popular around the globe. A major trend in the current studies on social multimedia is using the social media sites as a source of huge amount of labeled data for solving large scale computer science problems in computer vision, data mining and multimedia. In this paper, we take a new path to explore the global trends and sentiments that can be drawn by analyzing the sharing patterns of uploaded and downloaded social multimedia. In a sense, each time an image or video is uploaded or viewed, it constitutes an implicit vote for (or against) the subject of the image. This vote carries along with it a rich set of associated data including time and (often) location information. By aggregating such votes across millions of Internet users, we reveal the wisdom that is embedded in social multimedia sites for social science applications such as politics, economics, and marketing. We believe that our work opens a brand new arena for the multimedia research community with a potentially big impact on society and social sciences.

Categories and Subject Descriptors

H.4 [**Information Systems**]: Information Systems Applications; J.4 [**Computer Applications**]: Social and Behavioral Sciences

General Terms

Algorithms, Economics, Experimentation, Human Factors

MM'10, October 25–29, 2010, Firenze, Italy.

Keywords

Social multimedia, prediction, politics, economics, marketing

1. INTRODUCTION

Both governments and industries are interested in social trends. For example, politicians use polling to measure their popularity for elections and to monitor public sentiment to decide which position to take on social issues. Industry polls potential consumers to understand product acceptance. Although it is an expensive undertaking to perform polling, it is an investment that is critical for organizations both large and small to use for resource allocation and planning.

Figure 1: Flickr pictures of Obama, Hillary Clinton, Nokia phone, iPod, PS3 and Mac.

Social multimedia (including photos and videos) hosting and sharing websites, such as Flickr, Facebook, Youtube, Picasa, ImageShack and Photobucket, are gaining popularity around the world. A major trend in the current studies (such as Flickr distance[22] and tourism destination recommendation [4]) on social multimedia is using the social media sites as a source of huge amount of labeled data for solving large scale computer science problems in computer vision, data mining and multimedia. In this paper, we take a new path to explore the global trends and sentiments that can be drawn by analyzing the sharing patterns of uploaded and downloaded social multimedia. In a sense, each time an image or video is taken or viewed, it constitutes an implicit vote for (or against) the subject of the image. This vote carries along with it a rich set of associated data including time and (often) location information. By aggregating such votes across millions of Internet users, we reveal the wisdom that is embedded in social multimedia sites for social science applications such as politics, economics, and marketing. Figure 1 shows a few Flickr images that are useful for this study.

We use Flickr as the example platform in our study for two reasons: (1) *Popularity*: launched in February 2004 and required by Yahoo! in March 2005, Flickr is now one of the most popular social image storage and sharing websites (hosting more than 4 billion images as of October 2009), which enables users to annotate tags and make notes and comments, thus forming a large online social network and information network at the same time; and (2) *Availability*: unlike other social network websites, such as Facebook and MySpace, most Flickr data is publicly available and downloadable via its API.

Our paper reveals the power of Flickr through the following case studies in diverse domains:

(1) **Politics**. We monitor the progress of a US presidential election. Traditionally, the popularity of candidates is determined by conducting polls, which have limitations in sampling, are time consuming, and do not necessarily reflect the most recent changes in public opinions. Web data may contain the most up-to-date information that can reflect the popularity of the candidates and their support among populations. Our study on the 2008 United States presidential election shows that Flickr provides hints that indicate the winners of the party presidential primaries and the presidential election itself.

(2) **Economics**. When people buy consumer products, such as cars and cell phones, it is likely that some buyers will take photos and upload them to social sharing websites. The number of related photos uploaded online can reflects the number of products sold in the market. Our study shows that Flickr can predict the product sales (such as for iPod and iPhone) before the companies release the quarterly reports. This gives a way of estimating the company sale performance in real time and on a world-wide scale, potentially offering the advantage of helping investors make better and earlier decisions in the stock market.

(3) **Marketing**. We can monitor product distribution around the world over time. In addition to product sales, by considering the geographic information (especially the GPS location) of photos, we are able to monitor the spread and adoption of a product around the world, which can help the company exploit the growing popularity in different regions for better planning and management of manufacturing, marketing and distribution.

We believe that our work opens a brand new arena for the multimedia research community with a potentially big impact on society and the social science studies.

The rest of the paper is organized as follows: Section 2 discusses related work. Section 3 proposes Flickr features. Section 4 presents prediction models. Section 5 reports experimental results and discussion. Section 6 concludes the paper. Section 7 shows out vision for the future.

2. RELATED WORK

Recently, James Surowiecki published the book "The Wisdom of Crowds" [21], espousing the idea that under the right conditions, a crowd of non-experts can lead to decisions that are even smarter than the experts within the crowd. The conditions include independence of crowd members, decentralization, diversity and a means for aggregating the judgements of members. For a website with a large user base such as Flickr, all of these conditions are met.

Further, other work shows that the actions of individual Internet users, when properly pooled, can indicate macro trends. There are studies using Search Engine queries for influenza Internet surveillance [8], such as Google Trends [12], search advertisement click through [11], Yahoo search queries [18] and health website access logs [15]. Specifically in [12], Google search engine queries and data from the Centers for Disease Control (CDC) are used to find 45 specific search terms that are related to the percentage of influenze related physician visits. This model allows for monitoring influenza rates 1-2 weeks *ahead* of the CDC reports.

The problem of using general search engines is that the original query log is not publicly available and the queries trends may become noisy under the impact of news events. For example, as soon as a new product is announced by a major technology company, blogs will begin to report and speculate about the product. However, images of the product do not become widespread until the product is in the hands of the public. We build a framework to show that social media (specifically, uploaded images), and not only search terms, are also useful for representing the consensus of a crowd and for forecasting trends and public opinions.

To the best of our knowledge, this paper is the first one to reveal the wisdom of social multimedia sharing website and demonstrate that it is a useful information platform for social sciences.

3. FLICKR FEATURES

This section presents indicative features extracted from Flickr, including features based on meta-information and image visual relevance. We also discuss the background model.

3.1 Meta-Information Features

From the Flickr website, we denote $\mathcal{T}$ as the set of all terms (included in tags, titles, descriptions, etc.), $\mathcal{I}$ as the set of all images, $\mathcal{U}$ as the set of users and $\mathcal{D}$ as the set of dates at the day level (to achieve a multi-scale time analysis, we denote $\mathcal{M}$ as the set of dates at the month level, $\mathcal{Q}$ as the set of dates at the quarter level, and $\mathcal{Y}$ as the set of dates at the year level).

For any image $i \in \mathcal{I}$, denote $T(i) \subseteq \mathcal{T}$ as the set of tag annotations of image i, $D(i) \in \mathcal{D}$ as the day when the image was taken (similarly we have $M(i)$, $Q(i)$ and $Y(i)$ for the coarser levels of month, quarter, year, respectively), and $U(i) \in \mathcal{U}$ as the user of the image.

For any user $u \in \mathcal{U}$, denote $I(u) \subseteq \mathcal{I}$ as the set of images uploaded by user u, $I(u, d) \subseteq \mathcal{I}$ as the set of images uploaded by user u and taken on date $d \in \mathcal{D}$.

Given a query $q \subseteq \mathcal{T}$, we search Flickr to obtain $I(q)$, the set of relevant images. The meta information, such as user, date, title, description and tags, is also be extracted. Denote $I_T(q) \subseteq I(q)$ as the set of relevant images that are tagged with the query term(s), i.e., $I_T(q) = \{i | i \in \mathcal{I}, q \subseteq T(i)\}$.

We formulate the following **features/indicators** for use in building prediction models,

- Images per day (IPD) is the number of relevant images per day. With this indicator, every relevant image will serve as one count. More specifically, given a query $q \subseteq \mathcal{T}$ and a day $d \in \mathcal{D}$, we define

$$IPD(q, d) \equiv |I(q, d)| = |\{i | i \in I(q), D(i) = d\}| \quad (1)$$

Similarly, we define images per month (IPM), images per quarter (IPQ), and images per year (IPY). If we only consider relevant images that are tagged with the query, then we have a new indicator of tagged images per day (TIPD), which is defined as

$$TIPD(q,d) \equiv |I_T(q,d)| \quad (2)$$
$$= |\{i | i \in I_T(q), D(i) = d\}| \quad (3)$$

Similarly, we can define TIPM, TIPQ and TIPY over the time intervals of month, quarter, and year, respectively.

- Images with unique users per day (UPD) is the number of images with unique users per day. A user may take multiple photos of the same object and upload them to Flickr. To avoid this redundancy, multiple photos from the same user only count once with these features.

$$UPD(q,d) \equiv |I_U(q,d)| \quad (4)$$
$$= |\{u | u \in \mathcal{U}, I(u,d) \cap I(q,d) \neq \emptyset\}| \quad (5)$$

Similarly, we have the following time interval extensions: UPM, UPQ, UPY, TUPD, TUPM, TUPQ and TUPY.

- Images with first-time unique users per day (FUPD) is the number of images with first time unique users per day. A person may take multiple photos of the same object on different dates. UPD will count the same user multiple times if he/she has images on multiple dates. To further reduce redundancy, we only count the user with related images on the first date. More specifically, given a query q and date d, we denote the first date that user u took relevant image(s) as

$$D_{first}(u,q) = min\{d | d \in \mathcal{D}, I(u,d) \cap I(q,d) \neq \emptyset\} \quad (6)$$

then we can define FUPD as

$$FUPD(q,d) \equiv |I_F(q,d)| \quad (7)$$
$$= |\{u | u \in \mathcal{U}, D_{first}(u,q) = d\}| \quad (8)$$

By considering only tagged images over coarser time intervals, we have the following extensions: considering month or quarter or year levels (FUPM, FUPQ, FUPY); considering only tagged images, such as TFUPD, TFUPM, TFUPQ and TFUPY.

3.2 Image Visual Relevance

In addition to the Flickr features proposed in Section 3.1, we can also add features that are based on the visual content of the images themselves. Although presently our experimental results (Section 5) use only the aforementioned features, in the future we plan to investigate using visual relevance features to improve the results. $R(i,q)$ is the relevance score which indicates the relevance of image i to query q. Flickr returns a list of images for a query by the text information. However, text information is not always reliable; some returned images may not be truly relevant to the query. By also considering the visual information, we should be able to better estimate the total relevance.

The first step is to obtain a subset of ground truth images that are relevant to the query, which can be done by manual labeling. Then we can compute the visual similarity between images with those ground truth images and choose the highest score as the relevance score.

Image visual similarity can be estimated from low-level image content features [10], such as color histogram, texture, edge histogram, Color Correlogram [13], CEDD [5], GIST, shape and SIFT [16]. By representing images as points in a n-dimension feature space with either a single type of feature or a combination of multiple types of features, we can compute the image relevance by calculating the distance (such as L2 and χ^2 test statistic [3]) between the points.

Note that we can also use the image-rich information network formed in Flickr community to consider both link and visual information to reinforce link and visual based similarities to get further improvement on image relevancy estimation, similar to [14].

By considering visual relevance, the indicators will be modified as follows:

$$VIPD(q,d) = \sum_{i \in I(q,d)} R(i,q) \quad (9)$$

$$VUPD(q,d) = \sum_{u \in I_U(q,d)} max\{R(i,q) | i \in I(u,d) \cap I(q,d)\} \quad (10)$$

For example, VIPD indicates the number of images captured on a particular day d that are judged to be visually relevant to a query q by comparing each potentially matching image with a set of ground truth images. Similarly, we can extend other indicators to consider visual relevance.

Table 1 summarizes the Flickr features by the information used, including Image (I), User (U), Tag (T) and Visual (V).

Table 1: Flickr features. $\mathfrak{D} = \{D, M, Q, Y\}$

Image	User	Tag	Visual
IP$\mathfrak{D}$={IPD, IPM, IPQ, IPY}	UP$\mathfrak{D}$ FUP$\mathfrak{D}$	TIP$\mathfrak{D}$ TUP$\mathfrak{D}$,TFUP$\mathfrak{D}$	VIP$\mathfrak{D}$, VUP$\mathfrak{D}$, VFU$\mathfrak{D}$,...

3.3 Flickr Background Model

The popularity of Flickr website itself changes over time, and this must be considered when performing analysis. We represent this background model as F_B. Since the change in popularity over time is not readily available directly from the company, we need other sources to estimate it, such as Alexa[1], Google Trends[2] and general queries on Flickr. However, Alexa does not provide the complete trend history for public access. And both Alexa and Google Trends have the *news outliers* problem: they may be biased to news impact, because some news events may generate more clicks to the website.

Based on the above analysis, we do not directly use them as the Flickr background model. To produce a better estimate of the Flickr popularity trend, we perform general queries on Flickr, such as bike and flower, which are general terms and usually not used in news topics. We take the Google Trends index for bike as an example, as shown in Figure 2. The curve indicates that except for a seasonal component, the query is general and no news outliers exist.

Figure 3 shows the TUPM of "bike" and "flower" on Flickr. We can see seasonal patterns just like in Google Trends. The difference is that we can also see a growing trend that indicates the growing popularity of Flickr website itself, which is not specific to any particular general query.

[1] Alexa Internet, Inc. http://www.alexa.com/
[2] Google Trends. http://www.google.com/trends

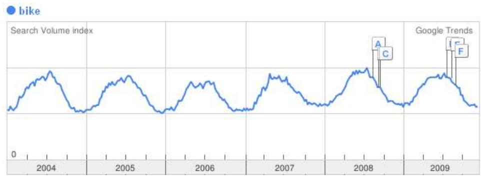

Figure 2: Google Trends search volume index for "bike".

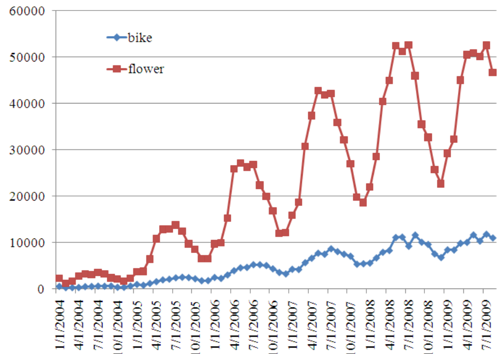

Figure 3: TUPM of "bike" and "flower" on Flickr.

In order to extract the background trend of Flickr, we use STL [6], provided in Package R [20], to decompose the TUPM time series for the general queries into Seasonal, Trend and Residual components. Figure 4 shows an example of the STL decomposition for query "bike". Query "flower" has similar results, although the scale is different.

We use several general queries instead of one to reduce bias. In order to estimate the true Flickr background trend, we first normalize each query's trend (because different queries have different scales) and then merge them together by computing the average value. More specifically, the Flickr background model F_B is computed as follows,

$$F_B = \frac{\sum_{i=1}^{m} Trend_i}{m} \tag{11}$$

where m is the number of general queries, and $Trend_i$ is the trend extracted by STL for query i.

4. PREDICTION MODELS

Forecasting is widely used in economy and marketing. Some models assume the time series data follows a stationary stochastic process, including autoregressive (AR) model and autoregressive moving average processes (ARMA) model. Some are based on a non-stationary process, such as Dickey-Fuller tests and long memory models. Those models have proved to be very compact and successful in predicting long term data. However, those models overlook the effects from users, and often fail to capture the development and evolution of new products or emerging events.

To forecast the sales of new products, Bass proposed to model the growth pattern by exploring the interaction between buyers and potential users [1] [2]. The model becomes one of the most famous empirical models in marketing, and is named as Bass model. Bass model implies sales growth to a peak and then decline, and provides a framework for guessing the long-term sales behavior based on early sales data.

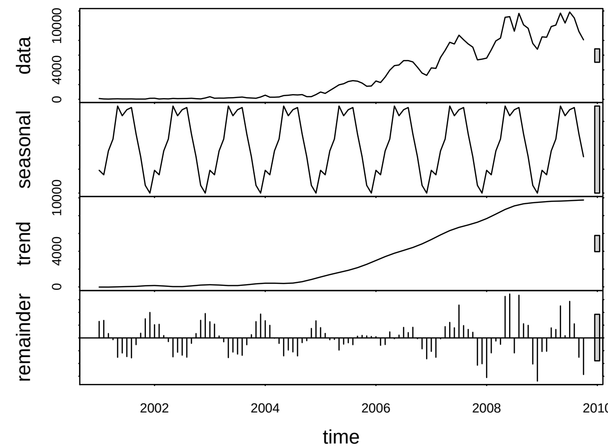

Figure 4: STL decomposition of TUPM time series for "bike" on Flickr.

As a successful example, in 1966 Bass accurately predicted that a sales peak would occur in 1968 for color television set sales at 5.7 million units, while the industry was (incorrectly) building plant capacity for 14 million picture tubes.

This section proposes new ways to predict product sales by the wisdom of social media. After providing our motivation, we briefly describe two popular economy models for forecasting product sales, autoregressive model and Bass diffusion model, and also present our extended models that consider the Flickr feature.

4.1 Motivation

In the new era of the Internet, user interests on products are strongly affected by everyday new information, and consequently exhibit complex patterns. A traditional way to study user interests and new technology acceptance is to use empirical surveys [7], where respondents are asked a set of structured questions and their responses are tabulated to capture the opinion of the respondents. However, survey methods often suffer from sampling bias and common methods bias [17]. To address such bias, researchers have to design survey questions carefully and validate data from multiple experiments, which makes a survey expensive and inefficient.

We propose to use the wisdom of crowd social media to estimate the user interests and conduct forecasting and predicting. Compared with traditional survey and prediction approaches, social media enjoys the following benefits:

- Social media information is inexpensive. Companies can reach a super-wide audience without worrying the cost of advertises or surveys. The nature of the social media allows consumers to exchange their interests with their friends or even strangers who share common interests. This information can be easily obtained by web crawlers.

- Social media has the advantage of measuring statistics almost instantaneously. Because exposure, response, and overall efficiency of Internet, social media are easier to track than traditional off-line approaches and can offer knowledge in real time through web data analysis.

We denote X_t ($t = \{1, 2, ..., n\}$) as the product sales at time t, and F_t as the Flickr feature value at time t. F_t can be any particular feature in Table 1. The goal is to predict future values of X_t.

4.2 Regression-based Models

Autoregressive (**AR**) model of order p, noted as $AR(p)$, is given by

$$X_t = c + \sum_{i=1}^{p} \alpha_i X_{t-i} + \varepsilon_t \tag{12}$$

where $\alpha_1, ..., \alpha_p$ are the parameters of the model, c is a constant and ε_t is the white noise.

Seasonal Autoregressive (**SAR**) model. Some time series may have seasonal factor, so we consider an extension of the $AR(p)$ model to $SAR(p, q, s)$ which includes the seasonal factor with order q. The seasonal model $SAR(p, q, s)$ is defined as

$$X_t = c + \sum_{i=1}^{p} \alpha_i X_{t-i} + \sum_{i=0}^{q} \beta_i X_{t-i-s} + \varepsilon_t \tag{13}$$

where q is the order used in the seasonal component, and s is the seasonal gap. Usually, when the date of the model is at month level, we set $s = 12$; for quarter level, $s = 4$.

Denote $\mathcal{F}_t$ as the Flickr index, which is defined as a Flickr feature adjusted by the background model,

$$\mathcal{F}_t = F_t + \omega F_B \tag{14}$$

where F_t is a Flickr feature, and F_B is the Flickr background model.

By considering the Flickr index, the AR model can be extended as

$$X_t = c + \sum_{i=1}^{p} \alpha_i X_{t-i} + \lambda\Theta\mathcal{F}_t + \varepsilon_t, \tag{15}$$

we refer to this model as **AR_Flickr**.

We define Θ as the scale factor which makes X_t and $\mathcal{F}_t$ in similar scale. It can be estimated from the training data,

$$\Theta = \frac{\sum_{t=1}^{n} X_t}{\sum_{t=1}^{n} \mathcal{F}_t} \tag{16}$$

Similarly, by adding the Flickr index, we can extend the seasonal SAR model to be

$$X_t = c + \sum_{i=1}^{p} \alpha_i X_{t-i} + \sum_{i=0}^{q} \beta_i X_{t-i-s} + \lambda\Theta\mathcal{F}_t + \varepsilon_t, \tag{17}$$

and we call this model as **SAR_Flickr**.

4.3 Diffusion-based Models

4.3.1 Bass Model

The Bass model [1] [2], developed by Frank Bass, is one of the most famous diffusion models with empirical generalizations in economics and marketing, and is widely used in product sale and technology forecasting. It models how a new product gets adopted as an interaction between users and potential consumers.

In the Bass model, the accumulative number of units ever sold at time t, $N_t = \sum_{i=1}^{t} X_t$, depends on three parameters: m, p and q. m is the total number of people who eventually buy the product. p is the coefficient of innovation. q is the coefficient of imitation.

The basic idea of the Bass model is that initial sales depends on the number of people who are interested in the novelty of the product, whereas later sales are impacted by the number of people who are drawn to the product after seeing their friends and acquaintances use it [9]. The Bass formula is

$$N_{t+1} = N_t + p(m - N_t) + qN_t(m - N_t)/m \tag{18}$$

Based on the Bass formula, X_t can be estimated as

$$X_t = m\frac{(p+q)^2 e^{-(p+q)t}}{p[1 + (q/p)e^{-(p+q)t}]^2} \tag{19}$$

4.3.2 Bass_Flickr Model

Bass model tries to characterize how a new product gets adopted as an interaction between users and potential consumers. We propose to take the influence from social media into account, which can also reflect the influence among people on social media website.

In order to consider Flickr index, we propose **Bass_Flickr model** as an extension of the Bass formula as follows,

$$N_{t+1} - N_t = p(m - N_t) + qN_t(m - N_t)/m + \lambda\Theta\mathcal{F}_{t+1} \tag{20}$$

which includes the factor of the Flickr index.

This formula is a difference equation and the solution is

$$N_t = m\frac{1 - e^{-(p+q)t}}{1 + (q/p)e^{-(p+q)t}} + \lambda\sum_{i=1}^{t}\Theta\mathcal{F}_i \tag{21}$$

Compute the pdf of the cumulative distribution function (cdf) based on Formula 21, and then we can derive the formula to compute X_t as follows,

$$X_t = \frac{d}{dt}N_t = m\frac{(p+q)^2 e^{-(p+q)t}}{p[1 + (q/p)e^{-(p+q)t}]^2} + \lambda\Theta\mathcal{F}_t \tag{22}$$

5. EXPERIMENTS

This section presents experimental results in three interesting areas to show the wisdom of Flickr: the 2008 U.S. Presidential election, monitoring product distribution around the world, and predicting product sales.

5.1 Presidential Election

The traditional way of predicting an election winner is by conducting polls. However, polls have limitations in sampling, are time consuming and do not necessarily reflect the most recent changes in public opinions. Web data can provide more instantaneous information that reflects the popularity of the candidates and their support among populations.

In a very real sense, each time an image of a politic candidate is taken, it constitutes an implicit vote for the candidate in the image. So the popularity of a candidate is reflected by the number of related images on Flickr.

This section presents our study on the 2008 United States presidential election. It shows that Flickr provides hints that indicate the winner of both Democratic and Republican party presidential primaries, and the general presidential election.

5.1.1 Democratic Party Presidential Primaries

Barack Obama, Hillary Clinton and John Edwards were major competitors during the 2008 Democratic Party presidential primaries. Obama announced his candidacy for president on 2/10/2007 and was named the presumptive nominee on 6/3/2008. We collect Flickr data from 1/1/2007 to 6/3/2008 with "Obama", "Hillary[3]" and "Edwards" as the query respectively.

Figure 5 shows the TUPD score of Obama, Clinton and Edwards on Flickr. We found that the peaks of the curves correspond to major campaign or voting events, such as Obama announcing his president campaign in Springfield (2/10/2007) [A], Obama winning the Iowa caucuses (1/3/2008) [B] and Super Tuesday (2/5/2008) when Obama won the majority of the available delegates [C].

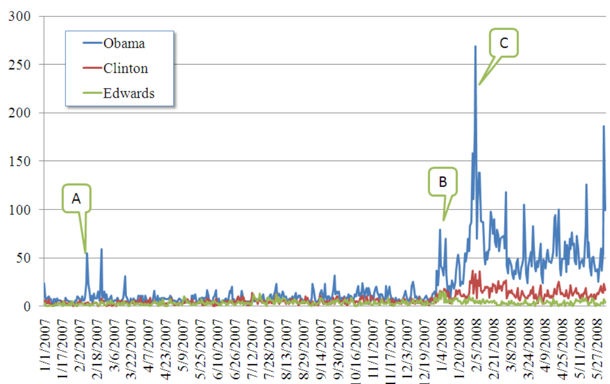

Figure 5: Democratic Primaries. TUPD for Obama, Clinton and Edwards.

Figure 6 shows the TUPM score of the three candidates, which gives a month level and compact overview of the general trends of their popularity. We can see that Obama was slightly more popular than Clinton and Edwards, but after winning Iowa, he gathered increasing momentum and continued to gain more popularity and finally won the nomination of the Democratic Party. Edwards was a major contender but then dropped out in January 2008; this is also indicted in the Flickr score.

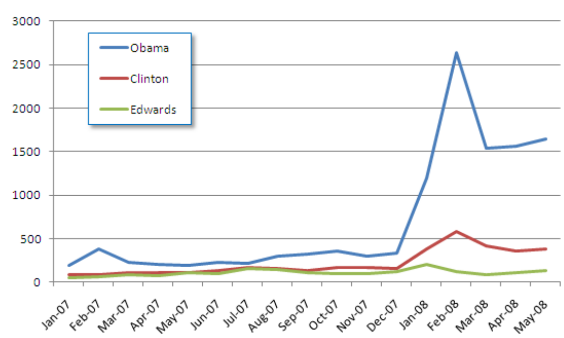

Figure 6: Democratic Primaries. TUPM for Obama, Clinton and Edwards.

Figure 7 shows an example of overlaying poll (conducted by Reuters/Zogby[4]) result together with the scaled Flickr index (using TUPM). We can see that the Flickr index closely matches the poll trend.

[3]In order to avoid images about Bill Clinton, we use Hillary as the query instead of Clinton

[4]http://www.pollingreport.com/wh08dem.htm

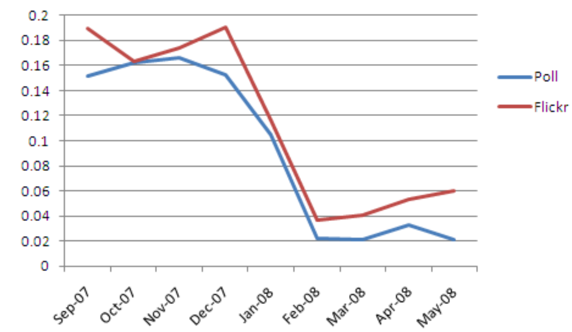

Figure 7: Reuters/Zogby Poll v.s. Flickr. Y-axis denotes the percentage of popularity for candidate Edwards.

5.1.2 *Republican Party Presidential Primaries*

McCain, Huckabee and Romney were the major competitors during the 2008 Republican Party presidential primaries. McCain ultimately won the nomination of the Republican Party. We collect Flickr data from 11/1/2007 to 6/1/2008 with "McCain", "Huckabee" and "Romney" as the queries respectively.

Figure 8 shows the TUPD scores of the three candidates. The winner of major primaries can be indicated by Flickr data, such as Huckabee winning the Iowa caucuses (1/3/2008) [A], McCain winning the New Hampshire primary (1/8/2008) [B], Romney winning the Michigan primary (1/15/2008) [C] and McCain winning Super Tuesday (2/5/2008) [D].

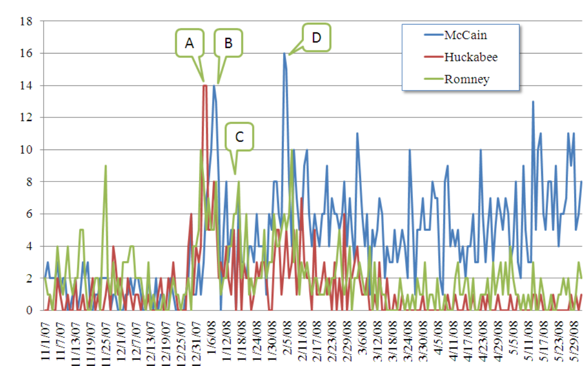

Figure 8: Republican Primaries, TUPD of "McCain", "Huckabee" and "Romney".

Figure 9 shows the TUPM score by percentage of the three candidates. It accurately reflects the candidates' popularity trends. At the beginning, Romney was widely considered as the major contender, then Huckabee won the first state voting in Iowa and briefly peaked in the polls, and finally McCain's popularity surged and he eventually won the Republican nomination.

5.1.3 *General Election*

During the 2008 US presidential general election, Democrat nominee Obama and Republican nominee McCain were the major competitors. Of course, Obama finally won the election

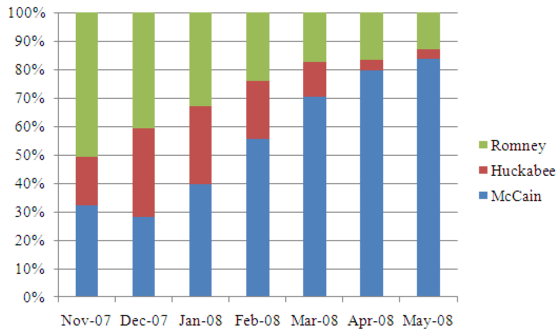

Figure 9: Republican Primaries. TUPM by percentage.

on 11/4/2008. We collect Flickr data from 6/1/2008 to 12/1/2008 with "Obama" and "McCain" as the queries, respectively.

Figure 10 shows the TUPD scores of the two candidates. Major events are reflected in the peaks of the TUPD curves, such as Democrat National Convention (8/25-28/2008) [A], Republican National Convention (9/1-4/2008) [B], first debate (9/26/2008) [C], second debate (10/7/2008) [D] and third debate (10/15/2008) [E].

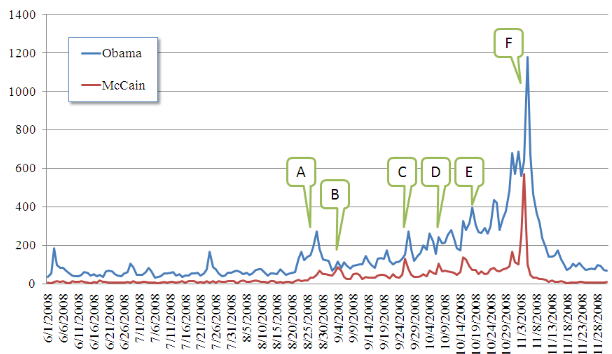

Figure 10: General election. TUPD of Obama and McCain.

An interesting finding: Obama won the president election on 11/4/2008 by capturing 52.9% of the popular vote [19]. The TUPD score for Obama and McCain on the same day is 642 (Obama) v.s. 571 (McCain) [F], with Obama got 642/(642 + 571) = 52.9% exactly the same as the real voting result. In addition, after Obama's win, a huge amount of press followed, and thus we see a high peak of TUPD on the next day.

Discussions. We may go deeper to analyze the geo information of the images and predict the winner in each state. We may also use the peaks to track campaign stops and the popularity of the corresponding candidate during each campaign location. In general, our study shows that social media correlates well with events and sentiments occurring in the political world.

5.2 Product Distribution in the World

By considering the geographic information (especially the GPS[5] location) of photos, we are able to monitor the spread and adoption of a product around the world over time, which can help the company (or its competitors) exploit the growing popularity in different regions for better planning and management of manufacturing, marketing and distribution.

We take the *iPod* as an example, which is a portable media player product first launched in October 2001 by Apple Inc.[6] and now very popular over the world. Figure 11 shows the geo-tagged iPod images on Flickr distributed over the world during years from 2006 to 2009, using mapping package M_Map[7]. Red points indicate new geo-tagged iPod images within the year compared with blue points indicating those appear before the year. The pattern shows how iPod spreads around the world over time. iPod was originally popular in North American and Europe, and then spread to other continents, such as Asian, Africa, South American and Australia.

5.3 Product Sales Prediction

When people buy products, such as cars and cell phones, it is likely that some will take photos and upload them to social sharing websites, as shown in Figure 1. The number of related photos uploaded online can reflect the number of product sales. Companies usually publish financial reports quarterly and some of them may release product sales data, which are available from the companies or some third-party websites, such as SECFilings[8].

In this section we show that Flickr features can provide successful prediction of product sales *before* companies release their quarterly reports. This provides a way of estimating a company's sales performance in real time and on a worldwide scale, and can potentially offer information for investors to make better and earlier decision in the stock market.

We conduct experiments on popular products, including music player (iPod), computers (Dell and Mac/Macintosh), cell phones (iPhone, Motorola[9], Nokia[10]), electric game console PS3[11]. The sales data for iPod, iPhone and Mac are available from Apple quarterly sales reports[12].

We focus on the task of predicting the quarterly product sales units. *Why quarterly sales?* Because usually only quarterly data are publicly available from the companies. *Why sale units instead of sale revenue?* We do not predict the sale revenue, because different products have different prices, and even the same product may have different prices at different times and different locations. We need the detailed price information sold for each unit to predict sale revenue, which is not available.

When building the prediction models, we use the quarterly Flickr feature TUPQ which is aggregated from the daily features. So both training and prediction are performed at the quarter level.

5.3.1 Evaluation Measures

To evaluate the prediction performance, we use the Abso-

[5]GPS. http://www.gps.gov/

[6]Apple Inc. http://www.apple.com/

[7]M_Map. http://www.eos.ubc.ca/~rich/map.html

[8]http://secfilings.com

[9]http://secfilings.com/SearchResults.aspx?ticker=MOT

[10]http://www.nokia.com/about-nokia/financials/quarterly-and-annual-information

[11]Sony quarterly consolidated financial results. http://www.sony.net/SonyInfo/IR/financial/fr/index.html

[12]Apple quarterly sales report. http://www.apple.com/pr/library

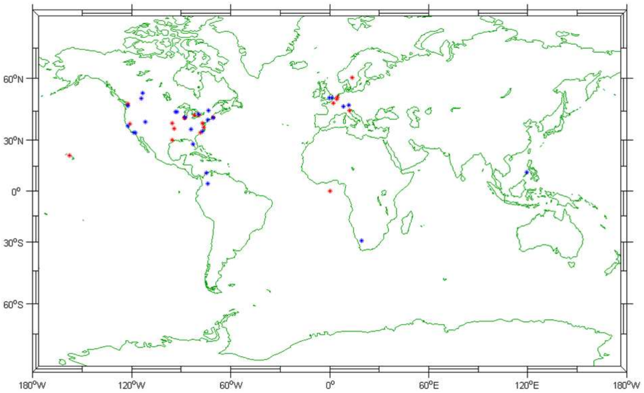

(a) 2006

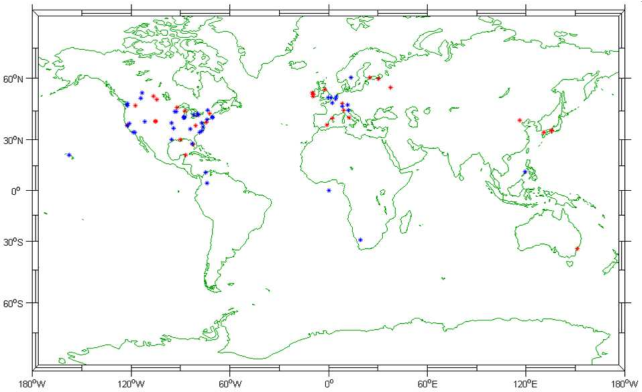

(b) 2007

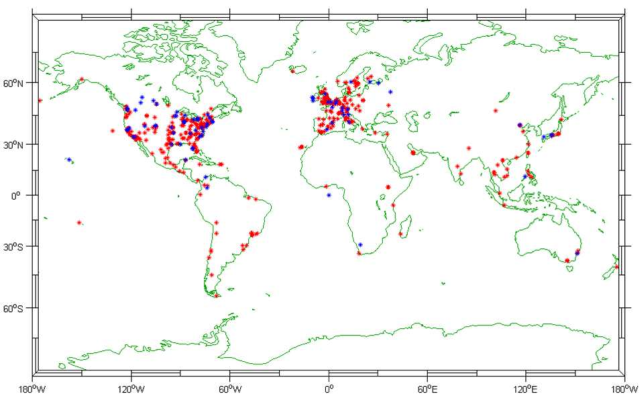

(c) 2008

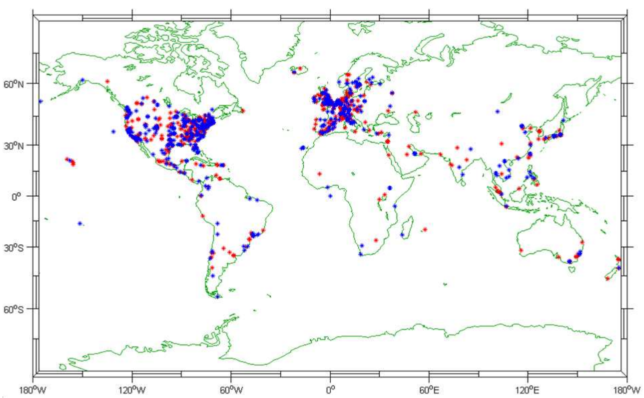

(d) 2009

Figure 11: *iPod* distribution around the world over the years. Red points indicate new geo-tagged iPod images within that year while blue points correspond to those before that year.

lute Error Ratio (AER),

$$AER = \frac{|P - S|}{S} \tag{23}$$

where P is predicted value, and S is the sales units. Smaller AER score indicates better prediction result. When $AER = 0$, the model achieves perfect prediction result.

The overall performance of a prediction model for m products is estimated by Mean Absolute Error Ratio (MAER),

$$MAER = \frac{AER_i}{m} \tag{24}$$

where AER_i is the AER score of the prediction for product i.

5.3.2 *Case Studies*

Figure 12 shows the iPod quarterly sales, Flickr features TUPM and TIPM. We can see that there is a strong correlation between the product sales and Flickr features.

Figure 13 shows the prediction results for iPod using model SAR_Flickr since it has a strong seasonal factor. Figure 14 shows the prediction results for Mac using model AR_Flickr. The red part is the predicted value for the last four quarters, and the green part is the model prediction for the training data.

We can see that our model not only fits the training data well, but also produces highly accurate predictions for future sales.

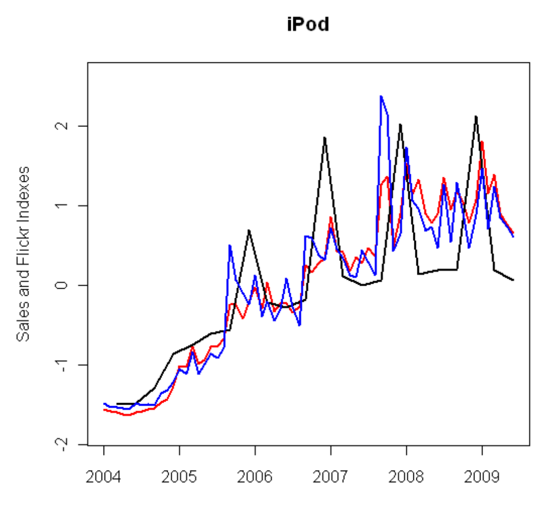

Figure 12: Normalized Sales v.s. Flickr trends for iPod. Sales (black), Flickr TUPM (red), Flickr TIPM (blue). Y-axis is the standardized value instead of the original value, because they are in different scales.

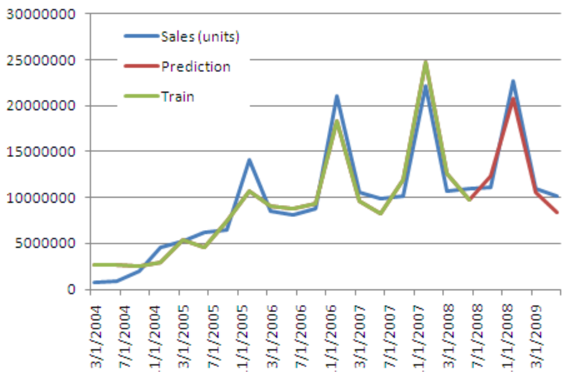

Figure 13: Prediction of iPod sales. Sales (blue), training fit (green), prediction (red). X-axis denotes sale units in the quarter level.

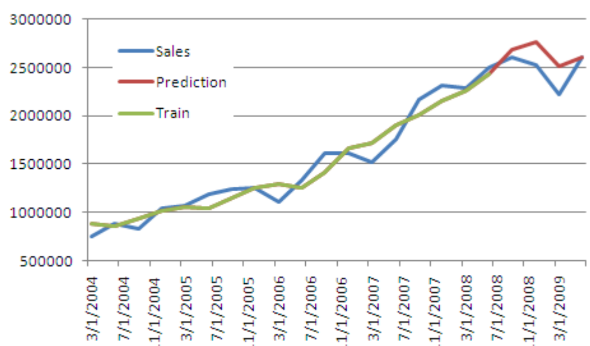

Figure 14: Prediction of Mac sales. Sales (blue), training fit (green), prediction (red). X-axis denotes sale units in the quarter level.

5.3.3 *Overall Performance*

Figures 15, 16 and 17 show the MAER score and the standard deviation of the models for predicting the sales of the products. We compare three pairs of models: AR v.s. AR_Flickr, SAR v.s. SAR_Flickr and Bass v.s. Bass_Flickr.

We can see that the models (AR_Flickr, SAR_Flickr and Bass_Flickr) extended to consider the Flickr features can achieve overall lower prediction error compared with traditional models (AR, SAR and Bass) which only consider sales history. In general, considering the Flickr features gives more robust results, with comparable standard deviation with the SAR models and much lower standard deviation than the AR and Bass models.

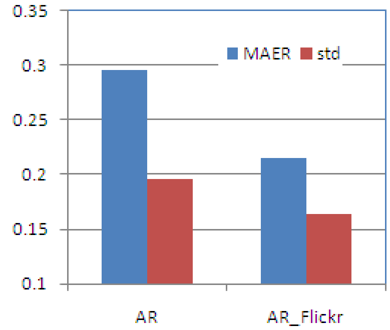

Figure 15: Overall performance of AR v.s. AR_Flickr. Y-axis denotes the MAER score and the std (standard deviation) among all the products.

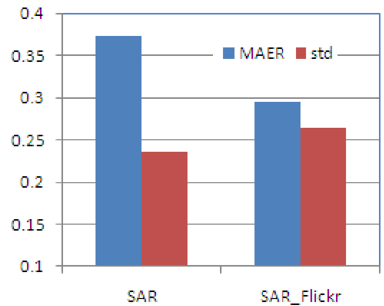

Figure 16: Overall performance of SAR v.s. SAR_Flickr. Y-axis denotes the MAER score and the std (standard deviation) among all the products.

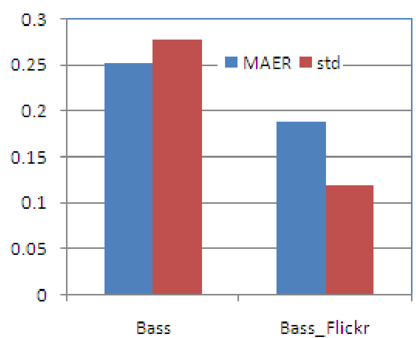

Figure 17: Overall performance of Bass v.s. Bass_Flickr. Y-axis denotes the MAER score and the std (standard deviation) among all the products.

6. CONCLUSIONS

We conclude this paper by summarizing our major contributions as follows:

(1) We show the wisdom of social multimedia and demonstrate its usefulness in serving as an information platform for social science studies, such as politics, economics and marketing.

(2) We propose several Flickr features considering images, tags and users for building prediction models and monitoring trends.

(3) We study on the 2008 United States presidential election and show that Flickr provides hints that indicate the winners of the party presidential primaries and the presidential election itself.

(4) Based on our Flickr features, we propose three new models (AR_Flickr, SAR_Flickr and Bass_Flickr) to extend widely used traditional prediction models in economics and marketing. Experiments on popular products show that our

models produce much better prediction performance both in terms of lower error rate and higher robustness.

(5) By considering the geographic information (especially the GPS location) of Flickr photos, we show that Flickr is able to monitor the spread and adoption of a product around the world. This can help the company exploit the growing popularity in different regions for better planning and management of manufacturing, marketing and distribution.

7. VISION FOR THE FUTURE

We believe this paper opens a brand new arena for the multimedia research community with a potentially big impact on society and social sciences.

Among many possible future work topics, we plan to investigate using visual relevance features to improve the results, build mathematic prediction models for presidential election winner prediction and utilize the geo-information of photos to track the candidate campaigns. In addition to product sales, it is also interesting to extend our models for predicting other popular economy metrics, such as GDP and unemployment rate.

There are some challenges, such as how to identify relevant images to the topic, how to handle missing or noise annotations, and how to integrate multiple social multimedia sources (we may even go further to include search engine query logs and any other online information) to get the most comprehensive information for better prediction.

Acknowledgment

Research was sponsored in part by Kodak Inc., NSF grants IIS-09-05215, AFOSR MURI award FA9550-08-1-0265, and by the Army Research Laboratory under Cooperative Agreement Number W911NF-09-2-0053 (NS-CTA). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

8. REFERENCES

[1] F. M. Bass. A dynamic model of market share and sales behavior. In *Proceedings of Winter Conference American Marketing Association*, Chicago, IL, 1963.

[2] F. M. Bass. A new product growth for model consumer durables. *Management Science*, (15):215–227, 1969.

[3] S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(4):509–522, 2002.

[4] L. Cao, J. Luo, A. Gallagher, X. Jin, J. Han, and T. S. Huang. A worldwide tourism recommendation system based on geotagged web photos. In *ICASSP*, 2010.

[5] S. Chatzichristofis and Y. Boutalis. CEDD: Color and edge directivity descriptor: A compact descriptor for image indexing and retrieval. *Lecture Notes in Computer Science*, pages 312–322, 2008.

[6] R. B. Cleveland, W. S. Cleveland, J. E. Mcrae, and I. Terpenning. STL: A seasonal-trend decomposition procedure based on loess. *Journal of Official Statistics*, 6(1):3–73, 1990.

[7] J. Converse. *Survey research in the United States: Roots and emergence 1890-1960*. Transaction Pub, 2009.

[8] C. D. Corley and A. R. Mikler. A computational framework to study public health epidemiology. In *International Joint Conferences on System Biology, Bioinformatics and Intelligent Computing (IJCBS09)*, Shanghai, China, August 2009.

[9] P. S. Cowpertwait and A. Metcalfe. *Introductory Time Series with R*. Springer-Verlag New York Inc., 2009.

[10] R. Datta, D. Joshi, J. Li, and J. Z. Wang. Image retrieval: Ideas, influences, and trends of the new age. *ACM Computing Surveys*, 40(2):1–60, April 2008.

[11] G. Eysenbach. Infodemiology: tracking flu-related searches on the web for syndromic surveillance. In *AMIA 2006 Symposium Proceedings*, pages 244–248, 2006.

[12] J. Ginsberg, M. H. Mohebbi, R. S. Patel, L. Brammer, M. S. Smolinski, and L. Brilliant. Detecting influenza epidemics using search engine query data. *Nature*, 457:1012–1014, 2009.

[13] J. Huang, S. R. Kumar, M. Mitra, W.-J. Zhu, and R. Zabih. Image indexing using color correlograms. In *Proceedings of the 1997 Conference on Computer Vision and Pattern Recognition (CVPR '97)*, page 762, Washington, DC, USA, 1997. IEEE Computer Society.

[14] X. Jin, J. Luo, J. Yu, G. Wang, D. Joshi, and J. Han. iRIN: image retrieval in image-rich information networks. In M. Rappa, P. Jones, J. Freire, and S. Chakrabarti, editors, *WWW*, pages 1261–1264. ACM, 2010.

[15] H. A. Johnsona, M. M. Wagnera, W. R. Hogana, W. Chapmana, R. T. Olszewskia, J. Dowlinga, and G. Barnas. Analysis of web access logs for surveillance of influenza. *Stud Health Technol Inform.*, 107(Pt 2):1202ÍC6, 2004.

[16] D. G. Lowe. Object recognition from local scale-invariant features. In *Proceedings of the International Conference on Computer Vision-Volume 2*, page 1150, Washington, DC, USA, 1999. IEEE Computer Society.

[17] P. Podsakoff, S. MacKenzie, J. Lee, and N. Podsakoff. Common method biases in behavioral research: A critical review of the literature and recommended remedies. *Journal of applied psychology*, 88(5):879–903, 2003.

[18] P. M. Polgreen, Y. Chen, D. M. Pennock, and F. D. Nelson. Using internet searches for influenza surveillance. *Clinical Infectious Diseases (Supplement)*, pages 1443–1448, 2008.

[19] R. C. Politics. General election: Mccain vs. obama. `http://www.realclearpolitics.com/epolls/2008/president/us/general_election_mccain_vs_obama-225.html`.

[20] R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2009. ISBN 3-900051-07-0.

[21] J. Surowiecki. *The Wisdom of Crowds*. Anchor, 2005.

[22] L. Wu, X.-S. Hua, N. Yu, W.-Y. Ma, and S. Li. Flickr distance. In *MM '08: Proceeding of the 16th ACM international conference on Multimedia*, pages 31–40, 2008.